



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 11 Issue: IV Month of publication: April 2023

DOI: <https://doi.org/10.22214/ijraset.2023.50583>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Modeling Radiologist Decision Making for Paediatric Elbow Fractures by Deep Learning Multiview

Aditya Kalyankar¹, Ganesh Aravind Harkude², Lekkala Sai Sumanth³

^{1,2}AI&ML BMSIT&M

³ECE BMSIT&M

Abstract: To detect an elbow fracture, patients typically need to take frontal and side views of diagnostic elbow radiographs. For the classification of elbow fracture subtypes, we propose a multiview deep learning technique in this work. Our strategy makes advantage of transfer learning by first training two individual models, one for the top view and another for the lateral position, and then moving the values to the relevant layers in the proposed multiview network design. Quantitative medical data has also been included into the training phase using a specific curriculum architecture that lets the model to initially learn from "easier" examples and then advance to "harder" examples to attain improved performance. Furthermore, our multiview network can work with both two simultaneous views as well as a single view as input. Using a database of 1,964 photographs for the classification of elbow fracture, we rigorously evaluate our methods. Results show that our technique may enhance the effectiveness of the tested methods and perform better than two comparable methods on broken bones investigation under various conditions.

The sampling probability of each training sample is determined by a scoring criterion that was built using clinically reliable knowledge from human specialists, where the scoring indicates the diagnostic complexity of different elbow fracture subtypes. Additionally, we provide a technique that modifies the sampling probabilities at each epoch and is applicable to many sample-based curriculum learning frameworks. We design an experiment utilising elbow X-ray images from 1865 for a binary task of fracture/normal and compare our recommended strategy to a base classifier and an earlier method using different criteria. Our data show that the recommended technique performs the best in terms of categorization. The effectiveness of the earlier method is also enhanced by our recommended probability update mechanism.

I. INTRODUCTION

Integrating information from many points of view is crucial for human cognitive ability. Patients typically require both the mirror image (also called as the front posterior view) and medial view of the forearm radiography to identify elbow fractures. This is due to the fact that specific fracture types may be easier to notice from a particular angle: the lateral view reveals the coronoid processes and the olecranon process, while the frontal aspect projects the proximal humerus, distal ulna, and the radius. In practise, it is also common for some individuals to only receive a single image radiography taken or to receive a viewpoint that is absent for several reasons.

Recent advancements in deep learning has made it feasible to use several views of X-ray images to automatically diagnose bone fractures [12,3,10]. When compared to utilising human experts, this process is faster and more precise. Multiview data, which provides extra visual information from different perspectives for detecting elbow fractures, is only utilized by a small number of approaches.

Using elbow radiographs collected from the front and lateral perspectives, respectively, we introduce a novel Multiview data augmentation network architecture in this research for categorizing elbow fracture subtypes. Given that it has a dual-view (frontal or lateral) architecture but doesn't necessarily require the dual input during inference, the proposed model is flexible. Our training method for the multiscreen model also takes advantage of transfer learning by first having trained two single-view designs, one for the full pic and another for the lateral aspect, and then transmitting the training parameters to the appropriate layers in the recommended multiscreen network architecture. Additionally, we investigate the advantages of including other medical perspectives in training using a curricular learning strategy that allows the model to initially learn from "easier" examples before going on to "harder" cases to enhance performance.

We offer a scoring system that, using clinically accepted knowledge from human professionals, indicates the diagnostic complexity of distinct elbow fracture subtypes. The curriculum is then built using permutations of the training data, sampling without replacements is carried out at the beginning of each epoch, and the sampling chance for each training sample is calculated using the scoring criterion. Our updating method modifies the sampling probabilities at each epoch and is extensible to other sample-based curriculum learning frameworks.



Figure No.1: In this investigation, an example of typical adult radiographs with AP, exterior, internal, and lateral projections were employed.

The standard random order. 6 Since that time, curriculum learning has been used to a variety of tasks, including image classification, object recognition, semantic segmentation, self- or semi-supervised learning, multi-task learning, and multi-modal learning (17–20). Even if human specialists are trained for years, few works make use of their outside expertise, which might be especially useful in the field of medical imaging.

We design an experiment employing a dataset of 1865 elbow X-rays from patients in order to evaluate our method by contrasting the findings with those from a baseline strategy and an earlier approach. The findings show that our strategy outperforms the methods that were tested, and the recommended probability update algorithm enhances the performance of the previous method.

II. RELATED WORK

Multiview learning [23] use data that has several viewpoints of the same objects. The term "co-training style algorithms" has replaced the earlier term "traditional Multiview learning techniques," which first emphasized semi-supervised training in which several views of data were successively added to the labeled set and educated by the classifier. Another class of Multiview learning techniques investigates Typically Given Learning (MKL), which was first developed to condense the search area of kernels. Recently, Multiview learning-based modelling has demonstrated encouraging outcomes, for instance in the detection of breast cancer and bone fractures.

Additionally, there is continuing research in curricular learning. Bengio et al. [1] first suggested it to enable machine learning to mimic human learning by first training a machine-learning model with "easier" data and afterwards moving on to "harder" instances. Research already done focuses on leveraging curriculum learning to incorporate domain knowledge into training. To incorporate domain knowledge, think about using the categorization complexity of several classes.

Study Conducted with MURA Dataset, This dataset is made accessible to the public for scholarly research through the "Bone X-Ray Deep Learning Competition" organized by the Stanford Machine Learning group. With a total of 40,561 images in png format, the elbow, finger, forearm, hand, femur, shoulder, and wrists all have X-ray images in MURA, one of the largest public radiographic images, categorized as either normal or pathological (fracture). In the classification study conducted by Rajpurkar et al. [13] using DenseNet-169 on this dataset, the overall Cohen's kappa score was 0.705 as well as the overall AUC score showing the area under the general Receiver Operator Characteristic (ROC) curve was 0.929 [2].

Classification study conducted on the shoulder Bone, Using the ResNet152 models on a total of 1891 cases, including 1376 cases of four different types of proximal humeral fracture, Chung et al. achieved a classification accuracy of 96% [15]. The Sezers CNN model classified 219 shoulder MR images in three groups with 98.43% accuracy: normal, oedema, and Hill-Sachs lesions [16].

The classification carried out by Urban et al. on 597 X-ray images of shoulders with implants showed the highest accuracy (80.4%) for the NASNet models that was pre-trained using ImageNet [17].

ResNet, DenseNet, VGG, and InceptionV3 are some of the CNN-based deep learning models created for categorization, as well as their modified versions with Spinal FC and proposed ensemble learning models. In Section 4, along with an elaboration of the open-source set of data, which includes the neck bone X-ray images, feature extraction, and marked procedures applied thereon, a table containing the precision, recall, F1-score, and Cohen's kappa results attained by classification models is provided. The study's addition to the body of knowledge is examined in the last section, along with any modifications that may be done for future studies.

III. METHODS

In this work, three-channel, png-formatted X-ray pictures of the shoulder bone were classified as normal or pathological using a variety of deep-learning-based techniques. First, classification was carried out using publicly accessible deep learning models based on CNNs called ResNet (34,50,101,152), ResNeXt (50,101), DenseNet (169,201), VGG (13,16,19), InceptionV3, and MobileNetV2. New classification networks were then created by substituting the SpinalNet FC layer for the classification layer in each of the models utilised in this article. To further improve the classification accuracy, new ensemble learning models tailored to this study were constructed based on the findings gained here. The next subsections include details on the ensemble learning models used for classification, freshly created models with SpinalNet, and built deep learning models.

A. Multiview Model Architecture

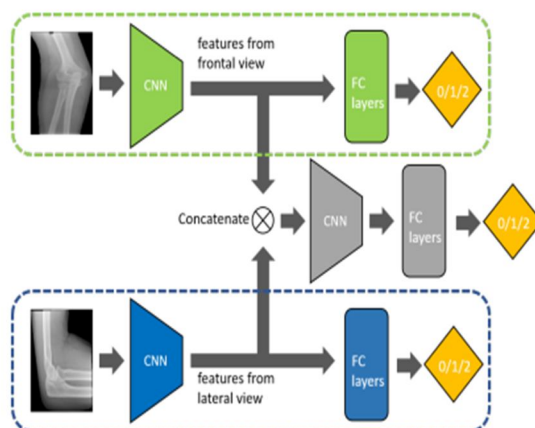


Figure No. 2: The Proposed Architecture

To combine data from both the frontal and lateral views of elbow X-ray images while maintaining the flexibility to be able to make predictions using just one view as input, we propose a special Multiview model architecture, which is shown in Figure 2. In this design, pairs of frontal and lateral view images are fed into the appropriate modules during training to allow convolutional neural networks to extract features (CNNs). Following feature extraction, Figure 2 shows how the model splits into three branches. Before performing classification, the centre branch further extracts features using the aggregated features from both perspectives as input. The fully connected (FC) layers get the matching single-view characteristics from the top and bottom branches, respectively, for classification.

Consider a sample data, $D_i = \{x_i, x_j, y_i\}$, where D_i represents the i th datasets, x_i, x_j are fractured images from training datasets, $y_i \in \{0,1,2\}$ is its ground truth label, with 0 representing no fracture, 1 ulnar fracture, and 2 radial fracture. The three projected labels from the three branches of our Multiview model are represented by the symbols $F(x^{(F)}_i)$, $L(x^{(L)}_i)$ and $M(x^{(F)}_i, x^{(L)}_i)$ where F, L, and M stand for the frontal view module (FVM), lateral view module (LVM), and "merge module" (MVM), which contains the two CNN blocks from the FVM and LVM, the CNN, and the FC layers in the middle branch.

The parameters in the overall model, the frontal view module, the lateral view module, and the merge module are represented by, F, L, and M during training. We minimize the objective function over the i -th data sample derived by equation (1). For each module, the loss is calculated using cross entropy loss over the associated predicted label and ground truth y_i in a one-hot form, as indicated in equation (1) (where C is the number of classes).

$$j_0 = (x^{(F)}_i, x^{(L)}_i, y_i) = j_{0f} = (x^{(F)}_i, y_i) + j_{0l} = (x^{(L)}_i, y_i) + j_{0m} = (x^{(F)}_i, x^{(L)}_i, y_i) = -\sum_{c=1}^C (\log(F(x^{(F)}_i)) + \log(L(x^{(L)}_i)) + \log(M(x^{(F)}_i, x^{(L)}_i)))$$

In the test phase, the merging module's final predicted label is used by default if both a frontal view picture and a lateral view image are supplied. Alternatively, even if there is just one view, the model will still provide a projected label from the module of that view, which will be acknowledged as a result of the model's specified architecture.

B. Learning Transfer from Previously Trained single-view Models

Researchers employ the ImageNet [5] pretrained model as a method of transfer learning in most deep learning applications in the field of medicine. Most deep learning models, particularly those that were self-designed, do not, however, contain publicly accessible pretrained weights. Here, we explore a homogeneous transfer learning method, as seen in Figure 3: We initially train two single-view models with the same frontal view and lateral view modules as the Multiview architecture (using the same training set as the one for the Multiview model).

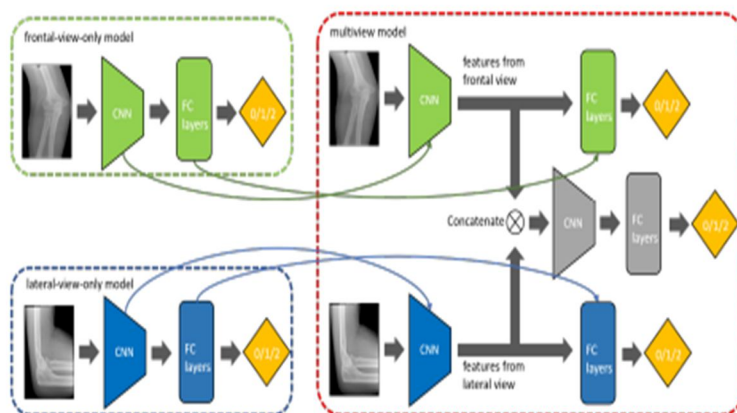


Figure No. 3: Learning transfer from previously trained single-view models

The learned weights of the CNNs and FC layers are then transferred from the single view models to their corresponding counterparts in the Multiview model (refer to the links in Figure 3). We initialize the weights of the middle branch's CNN and LC layers blocks in the merging module at random. In the Multiview model, we enable trainability for all weights.

C. Curriculum Learning Based on Knowledge

We suggest a knowledge-guided curriculum for model training to improve learning outcomes. The goal of curricular learning is to make it possible for the training process to proceed in a "easy-to-hard" order, with the easier examples being fed into the training model before the tougher ones. To achieve this, we modified the approach from [15] and built a Multiview-based curriculum. By rating the categorization complexity levels of each type of elbow fracture using board-certified radiologist's expertise, we quantify and integrate medical knowledge. The numeric values in Table 1 indicate the categorization difficulties based on the knowledge of experienced radiologists. To train the Multiview model, we utilize the "Both views" scores; however, for homogeneity transfer learning, we only use the "Frontal/Lateral view alone" scores.

	Normal	Ultra-fractured	Radial fractured
Front View	45	45	45
Lateral View	50	75	55
Both Views	55	80	65

Table No. 1 - The ability to integrate medical knowledge into curriculum learning is made possible by quantitative categorization difficulty levels for each type of elbow fracture (1-hardest to 100-easiest).

According to formula (2) with $e = 1$, these scores are used to activate the sampling probability for each training data point, where $p(i)$ represents the initial sampling probability for data point D_i , s_i represents its score, s_k represents the score of data point D_k , and N represents the total number of data points in the dataset. We permute the training set by sampling all the data points without replacement at the start of each epoch using the sampling probabilities.

$$p_i^{(e)} = \begin{cases} \frac{s_i}{\sum_{k=1}^N s_k} & e = 1, \\ p_i^{(e-1)} \cdot E' \sqrt{\frac{1/N}{p_i^{(0)}}} & 2 \leq e \leq E', \\ 1/N & E' < e \leq E \end{cases}$$

As a result, there is a greater possibility that the simpler samples will be offered before the more difficult ones. By changing the sample probabilities for each data point in accordance with equation, this possibility will be exponentially reduced (2). E_0 is the most recent epoch at which we updated the sampling probabilities, and e is the current epoch in this equation. The sample probabilities will be set to $1/N$ for the remaining training period ($E_0 \leq e \leq E$).

D. Training Set Generation

The remaining 20,350 instances served as the training dataset, of which 15 384 cases were determined to be positive for an acute or subacute aberration and 4966 cases to be negative. These reports did not analyze ambiguity. According to the instructions in Appendix E1, images were preprocessed and underwent data augmentation (supplement)

1) First Phase of Training

We divided our training into two phases. We altered the Xception (12) architecture to accept single-channel gray-scale input for phase 1's image classifier. Without using pre-trained weights, this model was trained. To distinguish it from the final model utilized in phase 2, we refer to this as the vision model (Fig. We changed the vision model's input size to $500 \times 500 \times 1$. (As represented by single-channel floating point values in the range of 0 to 1).

We added a two-unit dense layer following a scaled exponential linear units activation function after the final feature extraction layer of our model with adaptive global average pooling, where high-level image features are retrieved upon forward pass through the model (13). The nonlinearity required for our model's artificial neuron layers is provided by activation functions such as scaled exponential linear units, which are comparable to thresholding of synapses in biological neurons where a neuron only fires after meeting a particular level of activation (14).




Figure No. 4: The convolutional neural network was trained in phase 1 of the training process (CNN). Randomly selected pictures from the training set were given as single images, and the matching report from the related radiography series served as the ground truth for each image.

The amount of mistake the model produces is referred to as loss. The accuracy at the beginning of training is equal to random chance (50%), and the model has a large loss value. The loss of a neural network is often reduced by optimizers like Adam (name based on adaptive moment estimate) (15) by gradually lowering the network's weights (or parameters). Therefore, training as a whole may be seen as a process of loss optimization.

With the AMS Grad variation of the Adam optimizer, our model precisely optimized a cross entropy loss (16). Cross-entropy, which in our instance indicated the model's confidence that a certain picture is either positive or negative and whether the research was truly positive or negative, is a technique to characterize the inaccuracy between several categories and the ground truth (14).

The model was trained for 100 epochs, with the best model from the whole training run being utilized in phase 2. Each epoch was defined as one pass over the quantity of studies in the training set (20350 studies). In this investigation, epoch 83 had the lowest validation loss. If the loss did not decrease during 10 epochs to a minimum of 0.00001 (two decreases in total), the optimizer's learning rate, which was initially set to 0.001, was cut by 90%.

2) Second Phase of the Training

The trained CNN (visual model) served as a high-level feature extractor in a bigger model (Fig. 2) that included a recurrent neural network for the second round of training. During this stage of training, the CNN's parameters were locked. The 512-unit gated recurrent unit (GRU) and CNN were coupled via a 512-unit dense fully connected layer (17). A two-unit dense layer that produced floating point values came after the GRU. This model examined all three photos before making a choice after processing the first three in a sequence.

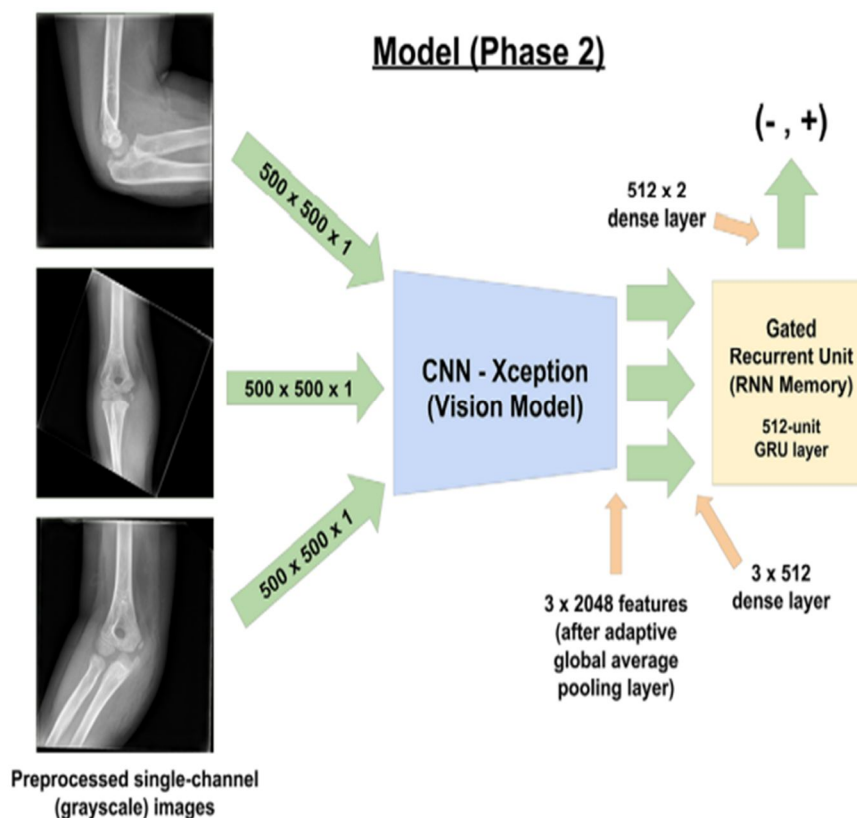


Figure No. 5 : The final model used to generate results from the validation set is represented by phase 2 of training. Prior to the final output of whether the study is positive or negative for abnormalities, the three radiographs can be processed in series using a recurrent neural network (RNN) and gated recurrent unit (GRU). The findings section contains a report and analysis of the model's outcomes. Convolutional neural network, or CNN.

We used the scores in the positive category to interpret the model's output values after they had been through a softmax layer (14) (Appendix E1 [supplement]). As a result, we utilised a straightforward scoring system where values between 0.0 and 0.5 were considered normal and values between 0.5 and 1.0 were considered aberrant. Low confidence values were those between 0.25 and 0.75, while high confidence values were those outside of this range. In Appendix E1 (extra), the distribution of output probabilities from the training and validation sets is shown for reference.

	Positive (+)	Negative (-)	
CNN (+)	536	82	PPV 87% (536/618) [95%CI: 84%, 89%]
CNN (-)	54	434	NPV 89% (434/488) [95%CI: 86%, 92%]
	Sensitivity 91% (536/590) [95%CI: 88%, 93%]	Specificity 84% (434/516) [95%CI: 81%, 87%]	Accuracy 88% (970/1106) [95%CI: 86%, 90%]

Table No. 2 : Pediatric elbow fractures are evaluated using a multiview convolutional neural network (CNN) in a contingency table.
NPV stands for "negative predictive value," and PPV for "positive predictive value."

3) Statistical Analysis

By calculating the sensitivity, specificity, positive predictive value, negative predictive value, area under the receiver operating characteristic curve (AUC), and their 95% confidence intervals, we assessed the performance of the classification algorithm on the validation set.




Figure No. 6 : The graph indicates that the model's area under the receiver operating characteristic (ROC) curve was 0.95 (95% confidence interval: (0.93, 0.960)).

IV. RESULTS

We contrast the curriculum learning technique (CL) and transfer learning (TL) used in our proposed Multiview model with the following six kinds of models: Single-view-frontal/lateral refers to two single-view models; Multiview refers to a multiscreen model with regular training; Multiview + TL refers to a multiview model with our proposed transfer learning strategy; Multiview + [11] refers to a multiview model with a previous curriculum training method; Single-view-frontal/lateral refers to a multiview model with two single-view models; Single-view + [11] + TL refers to a multiscreen The middle firm's output serves as the expected label in this sentence.

Our model has the flexibility to provide the prediction with a single view as input thanks to the various branches and the tailored loss function. We separately provide the performance data from the frontal view module and the lateral view module in Table 3. In contrast to [11], our curriculum updates each sample's difficulty score after every epoch, which is advantageous to the multiview model. Table 2 demonstrates that, compared to the state-of-the-art performance, our technique obtains the greatest AUC and balanced accuracy with a margin of up to 0.118.

False-negative fractures of the lateral and supracondylar condyles. These instances were the lone misses for each type of fracture in our test set. A, The field of view in the source anteroposterior view is too wide to clearly see the supracondylar fracture line. B, A zoomed and cropped version of the original image reveals the fracture line.

The positivity score is revised using the enlarged picture, rising from 0.14 (negative, high confidence) to 0.90. (positive, high confidence). The model is shown to be attentive to the healing fracture line in the lateral cortex (red arrow) and trabecular alterations in the medial supracondylar region in Saliency Map C (based on Guided Backpropagation of the Magnified Area) (blue arrow).




Figure No. 6 : Final Result

D, The field of view in the source anteroposterior view is too wide to clearly see the lateral condylar fracture line. E, The fracture line is visible when the image is cropped and zoomed to the region of interest. The positivity score is revised using the enlarged picture, rising from 0.44 (negative, low confidence) to 0.67. (positive, low confidence). F, The saliency map shows that the model is paying attention to the fracture line in the lateral condyle in the enlarged area (yellow arrow).

Application to Pediatric Musculoskeletal Radiology

No research that we are aware of have shown the effectiveness of deep learning models in paediatric musculoskeletal radiology. From a single frontal radiograph, Chung et al. showed a binary classification of adult proximal humeral fractures with a sensitivity and specificity of 99% and 97%, respectively (6). With a sensitivity and specificity of 90% and 88%, respectively, for adult patients, Kim and MacKinnon achieved distal radius fracture identification based just on a lateral image, and they eliminated any patient with open growth plates (7). In contrast to previous investigations, we sought to address a more complicated issue by avoiding a particular kind of fracture. Instead of having each image manually selected by a board-certified radiologist, we choose to filter our data using already-published radiology reports.




Figure No. 7 – False Negative Example

The former allows us to utilise a bigger dataset with the help of recently released natural language processing tools (11) to reduce classification mistakes, whereas the latter is both time-consuming and expensive. Furthermore, our dataset was more varied since it included data from individuals with both normal and aberrant radiographs at different ages and stages of skeletal growth. Our calculated mistake rate of 12% is within the range of previous reported error rates of radiologists' differences with physicians, peers, and second opinions radiologists, which vary from a low of 1.6% to a high of 41.8%. (27–29).

V. CONCLUSION

Our findings show that when there has been trauma, deep learning can efficiently binomially categorise acute and nonacute features on paediatric elbow radiographs. To the best of our knowledge, this is the first study that has effectively used deep learning to discriminate between open growth plates in skeletally immature patients and actual abnormalities. Our use of a recurrent neural network to classify a whole radiography series rather than just a single image and make a judgement based on all views, like a human radiologist would, is unique.

For the categorization of elbow fracture subtypes from frontal and lateral view X-ray images, we suggest an unique multiview deep learning approach. We pretrain two single-view models before using transfer learning on them. Through curricular learning, medical information was measured and included into the training process. The findings demonstrate that our Multiview model works better than the approaches that were tested, and we outperformed previously reported curriculum training techniques in terms of results. As part of our ongoing research, we intend to examine curricular learning in the output space and further incorporate additional domain knowledge considering various viewpoints.

REFERENCES

- [1] engio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: Proceedings of the 26th annual international conference on machine learning. pp. 41–48 (2009)
- [2] Cheng, C.T., Wang, Y., Chen, H.W., Hsiao, P.M., Yeh, C.N., Hsieh, C.H., Miao, S., Xiao, J., Liao, C.H., Lu, L.: A scalable physician-level deep learning algorithm detects universal trauma on pelvic radiographs. *Nature communications* 12(1),1–10 (2021)
- [3] Guan, B., Zhang, G., Yao, J., Wang, X., Wang, M.: Arm fracture detection in xrays based on improved deep convolutional neural network. *Computers & Electrical Engineering* 81, 106530 (2020)
- [4] Luo, J., Kitamura, G., Doganay, E., Arefan, D., Wu, S.: Medical knowledge-guided deep curriculum learning for elbow fracture diagnosis from x-ray images. In: *Medical Imaging 2021: Computer-Aided Diagnosis*. vol. 11597, p. 1159712. International Society for Optics and Photonics (2021)
- [5] Kingma DP, Ba J. Adam: A method for stochastic optimization. *arXiv [cs.LG] [preprint]* <https://arxiv.org/abs/1412.6980>. Published 2014. Accessed June 2018.
- [6] Reddi SJ, Kale S, Kumar S. On the convergence of Adam and beyond. In: *International Conference on Learning Representations*. http://www.sanjkv.com/AdamConvergence_ICLR.pdf. Published 2018. Accessed June 2018.
- [7] pringenberg JT, Dosovitskiy A, Brox T, Riedmiller M. Striving for simplicity: all convolutional net. *arXiv [cs.LG] [preprint]* <https://arxiv.org/abs/1412.6806>. Posted December 14, 2014. Revised April 13, 2015. Accessed June 2018.
- [8] Johnson R, Zhang T. Supervised and semi-supervised text categorization using LSTM for region embeddings. *arXiv [stat.ML] [preprint]* <https://arxiv.org/abs/1602.02373>. Posted February 7, 2016. Revised May 26, 2016. Accessed June 2018.
- [9] enugopalan S, Hendricks LA, Mooney R, Saenko K. Improving LSTM based video description with linguistic knowledge mined from text. *arXiv [cs.CL] [preprint]* <https://arxiv.org/abs/1604.01729>. Posted April 6, 2016. Revised November 29, 2016. Accessed June 2018.
- [10] Wei, J., Suriawinata, A., Ren, B., Liu, X., Lisovsky, M., Vaickus, L., Brown, C., Baker, M., Nasir-Moin, M., Tomita, N., et al., “Learn like a pathologist: curriculum learning by annotator agreement for histopathology image classification,” in [Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision], 2473–2483 (2021).
- [11] hang, D., Han, J., Zhao, L., and Meng, D., “Leveraging prior-knowledge for weakly supervised object detection under a collaborative self-paced curriculum learning framework,” *International Journal of Computer Vision* 127(4), 363–380 (2019).
- [12] Kervadec, H., Dolz, J., Granger, E., and Ayed, I. B., “Curriculum semi-supervised segmentation,” in [International Conference on Medical Image Computing and Computer-Assisted Intervention], 568–576, Springer (2019).
- [13] Wang, C., Zhang, Q., Huang, C., Liu, W., and Wang, X., “Manes: A multi-task attentional network with curriculum sampling for person re-identification,” in [Proceedings of the European Conference on Computer Vision (ECCV)], 365–381 (2018).
- [14] Tanzi, L., Vezzetti, E., Moreno, R., and Moos, S., “X-ray bone fracture classification using deep learning: a baseline for designing a reliable approach,” *Applied Sciences* 10(4), 1507 (2020).
- [15] Klein EJ, Koenig M, Diekema DS, Winters W. Discordant radiograph interpretation between emergency physicians and radiologists in a pediatric emergency department. *Pediatr Emerg Care* 1999;15(4):245–248.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)