

Medical Knowledge-Guided Deep Curriculum Learning for Elbow Fracture Diagnosis from X-Ray Images

Jun Luo^a, Gene Kitamura^b, Emine Doganay^b, Dooman Arefan^b, and Shandong Wu^{a, b, c, d}

^aIntelligent Systems Program, University of Pittsburgh, 4200 Fifth Avenue, Pittsburgh, PA, USA, 15213

^bDept. of Radiology, University of Pittsburgh, 4200 Fifth Avenue, Pittsburgh, PA, USA, 15213

^cDept. of Biomedical Informatics, University of Pittsburgh, 4200 Fifth Avenue, Pittsburgh, PA, USA, 15213

^dDept. of Bioengineering, University of Pittsburgh, 4200 Fifth Avenue, Pittsburgh, PA, USA, 15213

ABSTRACT

Elbow fractures are one of the most common fracture types. Diagnoses on elbow fractures often need the help of radiographic imaging to be read and analyzed by a specialized radiologist with years of training. Thanks to the recent advances of deep learning, a model that can classify and detect different types of bone fractures needs only hours of training and has shown promising results. However, most existing deep learning models are purely data-driven, lacking incorporation of known domain knowledge from human experts. In this work, we propose a novel deep learning method to diagnose elbow fracture from elbow X-ray images by integrating domain-specific medical knowledge into a curriculum learning framework. In our method, the training data are permuted by sampling without replacement at the beginning of each training epoch. The sampling probability of each training sample is guided by a scoring criterion constructed based on clinically known knowledge from human experts, where the scoring indicates the diagnosis difficulty of different elbow fracture subtypes. We also propose an algorithm that updates the sampling probabilities at each epoch, which is applicable to other sampling-based curriculum learning frameworks. We design an experiment with 1865 elbow X-ray images for a fracture/normal binary classification task and compare our proposed method to a baseline method and a previous method using multiple metrics. Our results show that the proposed method achieves the highest classification performance. Also, our proposed probability update algorithm boosts the performance of the previous method.

Keywords: Deep learning, curriculum learning, elbow fracture, medical knowledge

1. INTRODUCTION

Elbow fracture is one of the fracture types that happens most frequently among people across all ages. It is crucial for patients with elbow fracture to have timely diagnosis and treatment since the fracture could cause neurovascular damage.¹ Radiographic imaging, such as X-ray, is one of the most popular image modalities to help provide physicians with prompt assessment of the patient since it can help visualize the interior of the patient's elbow in a timely manner. Traditionally, it takes a physician years of training in order to diagnose elbow fractures from radiographic images. In recent years, thanks to the thrive of deep learning, a model that can classify and detect of different types of bone fractures can be trained within hours from annotated images to help diagnose elbow fractures.²⁻⁵ However, most existing methods purely rely on the data-driven learning, lacking leveraging known medical knowledge that physicians acquired after years of training.

In this work, we propose a novel deep learning framework that incorporates medical knowledge from domain experts (i.e., radiologists) through a curriculum learning⁶ method to augment the data-driven learning on a binary (fracture vs. normal) classification of the elbow fractures. Curriculum learning was first introduced to allow the machine to mimic human learning and train a model by feeding examples in an easy-to-hard order, opposite to

Correspondence to S. Wu. E-mail: wus3@upmc.edu, Telephone: 1 412 641 2567

This work is licensed under a [Creative Commons "Attribution 4.0 International"](https://creativecommons.org/licenses/by/4.0/) license. 




Figure 1. Six Subtypes of elbow fractures: (a) Ulnar fracture; (b) Radial fracture; (c) Humeral fracture; (d) Dislocation; (e) Complex fracture/multi-type fracture; (f) Coronoid process fracture.

the generally used random order.⁶ Since then, curriculum learning has been used in many areas such as image classification,^{5,7-10} object detection,^{11,12} semantic segmentation,^{13,14} self/semi supervised learning,^{7,11,12,14-16} multi-task learning,¹⁷⁻²⁰ multi-modal learning,^{7,8,21} etc.²²⁻²⁴ However, few works take advantage of the outside knowledge from human experts, which could be particularly helpful in the field of medical imaging since human experts are trained for years.

In our proposed framework, we design a scoring criterion, constructed based on clinically known knowledge from human experts, that indicates the diagnosis difficultness of different elbow fracture subtypes. The curriculum is then created through permutations of the training data by sampling without replacement at the beginning of each epoch and the scoring criterion is used to guide the sampling probability of each training sample. We also propose an algorithm that updates the sampling probabilities at each epoch, and this update algorithm is applicable to other sampling-based curriculum learning frameworks.

To evaluate our method, we design an experiment with an elbow X-ray dataset of 1865 patients and compare the effects with a baseline method and a previous method. Results show that our method outperforms the compared methods, and our proposed probability update algorithm boosts the performance of the previous method.

The remainder of this paper is organized as follows: Section 2 provides the details of the study cohort and the proposed method; Results of the experiment are shown in Section 3; Section 4 discusses our findings and limitations and concludes the study.

2. METHODS

2.1 Study cohort and dataset

Table 1. Number of images of the normal cases and six subtypes of the elbow fractures.

Type	(<i>normal</i>)	(<i>a</i>)	(<i>b</i>)	(<i>c</i>)	(<i>d</i>)	(<i>e</i>)	(<i>f</i>)	Total of (<i>a</i>)-(<i>f</i>) (<i>fracture</i>)	Total
Train	800	88	340	84	11	42	27	592	1,392
Test	400	10	44	9	2	4	4	73	473
Total	1,200	98	384	93	13	46	31	665	1,865

In this study, we use a cohort of 1,865 frontal view elbow radiographic images of patients with elbow trauma, where 665 are fracture positive cases and 1,200 are non-fracture (i.e., normal) cases. The elbow radiographic images were retrospectively collected at our institution and reviewed by expert radiologists. Among the 665 fracture cases, each case is annotated with a specific subtype of fracture out of a total of six different subtypes, as demonstrated in Figure 1. We focus on a binary classification task by grouping all six different subtypes of fractures as class fracture and assigning the non-fracture cases as class normal. We split the dataset into training set and test set while preserving a consistent percentage of each fracture subtype in the training and test set. Table 1 shows the distribution of our image dataset. During testing of the models, we evaluate effects on 100 randomly selected normal cases from the test set and all 73 fracture cases from the test set that serve as a more

balanced test between the binary classes, and we repeat the test 5 times, each time selecting different set of normal cases from the test set, to increase the robustness of the evaluation and report the average experiment results. We hereafter use the following notations to denote the six subtypes of elbow fracture: (a) for ulnar fracture; (b) for radial fracture; (c) for humeral fracture; (d) for dislocation; (e) for complex fracture/multi-type fracture; (f) for coronoid process fracture.

2.2 Knowledge-based curriculum

The idea of curriculum learning is first brought up to let the machine imitate human learning by first learning with “easier” samples, and then transit to learning from “harder” samples in the entire training set. In the process of constructing a curriculum, we need to design a quantifiable criterion that reflects how hard it is to classify a certain subtype of elbow fracture in clinical practice, and then utilize the criterion to schedule the order of feeding the samples into the model for training. Here, we use radiologists’ accumulated clinical experience on the difficultness of diagnosis of the elbow subtype fractures as a medical knowledge to formalize the criterion. As shown in Table 2, a board-certified experienced musculoskeletal radiologist (G.K.; > 10 years of experience) was asked to score the general level of difficulty of the elbow fracture diagnosis from the frontal radiographic view images of the elbow, with the scores ranging from 1 (hardest) to 100 (easiest).

Table 2. Number of images of the normal cases and six subtypes of the elbow fractures.

	<i>(normal)</i>	<i>(a)</i>	<i>(b)</i>	<i>(c)</i>	<i>(d)</i>	<i>(e)</i>	<i>(f)</i>
Score	30	30	30	70	40	90	10

In the training process, consider a triplet $(x_i, y_i, f_i) \in \{(x_j, y_j, f_j)\}_{j=1}^N$ in the training set, where x_i is the i -th image in the training set \mathcal{X} that contains N images, $y_i \in \mathcal{Y} = \{0, 1\}$ is its label, and $f_i \in \mathcal{F} = \{(normal), (a), (b), (c), (d), (e), (f)\}$ is its fine-grained label, indicating which subtype of the elbow fracture is, if it is not *(normal)*. As mentioned above, we focus on a fracture vs. normal binary classification task and $y_i = 0$ if and only if $f_i = (normal)$.

We design our curriculum with a sampling-without-replacement strategy to permute and reorder the training set at the beginning of each training epoch. For sampling, each x_i is given a probability, p_i , whose value is updated at the beginning of each training epoch, with the initial value $p_{i,(1)}$ assigned at the first epoch. Note that the probability update method mentioned below can be utilized by any curriculums learning framework that is based on permuting the training set by the strategy of sampling-without-replacement.

Initialization of p_i The $p_{i,(1)}$ for all images are computed from the hardness scores as provided in Table 2. Our curriculum addresses the importance of learning from easier cases, or, here, cases with higher scores by defining the initial probability $p_{i,(1)}$ for image x_i according to Equation (1), where s_{f_k} is the score for image x_k in terms of its fine-grained class label f_k , and N is the total number of images in the training set.

$$p_{i,(1)} = \frac{s_{f_i}}{\sum_{j=1}^N s_{f_j}} \quad (1)$$

Update of p_i At the beginning of each epoch after the first, the p_i ’s are updated by Equations (2) and (3):

$$\lambda_i = \sqrt[L]{\frac{p_{i,(final)}}{p_{i,(1)}}} = \sqrt[L]{\frac{1/N}{p_{i,(1)}}} \quad (2)$$

$$p_{i,(e)} = \begin{cases} p_{i,(e-1)} \cdot \lambda_i, & 2 \leq e \leq L \\ 1/N, & L < e \leq E \end{cases}, \quad (3)$$

where $p_{i,(e)}$ is the probability for image x_i after the update at epoch e , and λ_i is a scalar for x_i that ensures p_i for every image to exponentially and smoothly transit to $p_{i,(final)} = 1/N$. E denotes the total number of epochs,

and L is the number of the epoch after which p_i equals to $p_{i,(final)}$ for the rest of the training process. We treat L as a hyperparameter. In this way, the model can gradually reduce the emphases we initially put on the easier cases, and then reach to a stage of random selection of samples in the entire training set after epoch L .

3. RESULTS

We evaluate the performance of our curriculum on a total of 1,865 X-ray images (each image from a different patient). For each hyperparameter setting, we train the model 5 times to reduce potential bias from the randomness of sampling and training initialization. Results reported below are means and standard deviations across the 5 runs. All experiments are implemented in PyTorch²⁵ framework on an NVIDIA TESLA V100 GPU from Pittsburgh Supercomputing Center, with the same effort for hyperparameter tuning and data augmentation. We compare our method with two other methods: 1) a baseline method, which is the most commonly used method of random shuffling the order of the images in the training set; and 2) a previously reported curriculum learning method (here referred as MBDCL⁵) that uses a different way to permute the training set. In addition, in order to evaluate the effects under the opposite strategy of our curriculum, we also compare our method to an anti-curriculum setting where we reverse the difficulty scoring (listed in Table 2). We reverse the difficulty scoring by re-assigning the scores as 100 minus the original scores, indicating difficulty information that is the opposite to the knowledge provided by the radiologist. Furthermore, we evaluate the effects of plugging our probability update algorithm to the previous method (MBDCL + proposed update algorithm) and the anti-curriculum method (Anti-curriculum + proposed update algorithm). To ensure a fair comparison, all the methods use the same backbone of VGG16²⁶ network architecture. The evaluation metrics include accuracy, area under the ROC curve (AUC), Average Precision, and F1-score. Our experiment results are shown in Table 3. The ROC curves for the entire test set are shown in Figure 2.

Table 3. Performance comparison of different methods.

	Average on 5 Different Test Subsets (100 normal + 73 fracture)			
	Accuracy	AUC	Average Precision	F1 score
Baseline	0.776 ± 0.026	0.834 ± 0.025	0.788 ± 0.043	0.716 ± 0.032
MBDCL ⁵	0.797 ± 0.025	0.865 ± 0.019	0.831 ± 0.029	0.763 ± 0.022
MBDCL ⁵ + proposed update algorithm	0.806 ± 0.027	0.878 ± 0.022	0.847 ± 0.035	0.762 ± 0.029
Ours	0.809 ± 0.023	0.882 ± 0.020	0.852 ± 0.036	0.778 ± 0.024
Anti-curriculum	0.765 ± 0.032	0.863 ± 0.020	0.821 ± 0.036	0.757 ± 0.025
Anti-curriculum + proposed update algorithm	0.803 ± 0.036	0.871 ± 0.025	0.838 ± 0.045	0.773 ± 0.031

4. DISCUSSION AND CONCLUSION

In this work, we propose a novel medical knowledge-guided deep curriculum learning method for elbow fracture diagnoses from X-ray images. The knowledge is incorporated into the modeling process in the form of a pre-defined quantitative scoring criterion on the difficulty of classifying the elbow fracture subtypes. Various medical knowledge exists or can be accumulated/acquired from clinical experience. It will be a more effective way to augment the pure data-driven deep learning by leveraging the medical knowledge to the models, and to further improve the performance of deep learning. In this study, we formulate a technical framework to define and incorporate radiologists' diagnosis knowledge and we show in the results the following findings: first, our method outperforms all other compared methods (accuracy=0.81, AUC=0.88); second, the anti-curriculum settings demonstrate inferior results as expected, which reflects the positive impact of our original curriculum; in addition, we also show that the proposed probability update algorithm can further enhance other curriculum learning methods.

Despite the promising results, our work has some limitations. In our study, we are currently using a cohort of around 1900 images collected at our institution, where only around a third of which are fracture cases. A more




Figure 2. ROC curves of different methods for the entire test set.

balanced dataset and potentially multi-center larger cohorts would help further evaluate our method especially the robustness and generalizability.

There also exists some interesting topics for future investigation. Patients with elbow trauma are often suggested to take radiography from multiple perspectives (frontal and lateral views). Although plugging our method directly into a multiview deep neural network structure is feasible, more informative domain knowledge regarding different views could further be involved in the curriculum. We believe that the patterns and features learned by the model from different individual views could complement one another and possibly further improve the performance. In addition, based on different forms of clinical knowledge, customized schemes of combining multiple sources of knowledge are also worth investigation in future work.

In summary, in this work, we propose a novel strategy for incorporating medical knowledge to guide and enhance data-driven deep learning for medical applications especially for elbow fracture diagnosis. Our method demonstrates a new mechanism of defining and incorporating existing rich clinical experience and knowledge into artificial intelligence (AI) tools for clinical applications. We believe knowledge-guided AI represents a direction of the future research for AI in medical applications.

ACKNOWLEDGMENTS

This work was supported by National Institutes of Health (NIH)/National Cancer Institute (NCI) R01 grants (#1R01CA193603, #3R01CA193603-03S1, and #1R01CA218405), and an Amazon AWS Machine Learning Research Award. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation (NSF) grant number ACI-1548562. Specifically, it used the Bridges-2 system, which is supported by NSF award number ACI-1445606, at the Pittsburgh Supercomputing Center (PSC).

REFERENCES

- [1] Saeed, W. and Waseem, M., “Elbow fractures overview,” (2017).

- [2] Rayan, J. C., Reddy, N., Kan, J. H., Zhang, W., and Annapragada, A., “Binomial classification of pediatric elbow fractures using a deep learning multiview approach emulating radiologist decision making,” *Radiology: Artificial Intelligence* **1**(1), e180015 (2019).
- [3] England, J. R., Gross, J. S., White, E. A., Patel, D. B., England, J. T., and Cheng, P. M., “Detection of traumatic pediatric elbow joint effusion using a deep convolutional neural network,” *American Journal of Roentgenology* **211**(6), 1361–1368 (2018).
- [4] 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).
- [5] Jiménez-Sánchez, A., Mateus, D., Kirchhoff, S., Kirchhoff, C., Biberthaler, P., Navab, N., Ballester, M. A. G., and Piella, G., “Medical-based deep curriculum learning for improved fracture classification,” in [*International Conference on Medical Image Computing and Computer-Assisted Intervention*], 694–702, Springer (2019).
- [6] Bengio, Y., Louradour, J., Collobert, R., and Weston, J., “Curriculum learning,” in [*Proceedings of the 26th annual international conference on machine learning*], 41–48 (2009).
- [7] Gong, C., Tao, D., Maybank, S. J., Liu, W., Kang, G., and Yang, J., “Multi-modal curriculum learning for semi-supervised image classification,” *IEEE Transactions on Image Processing* **25**(7), 3249–3260 (2016).
- [8] Lotter, W., Sorensen, G., and Cox, D., “A multi-scale cnn and curriculum learning strategy for mammogram classification,” in [*Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*], 169–177, Springer (2017).
- [9] Wang, Y., Gan, W., Yang, J., Wu, W., and Yan, J., “Dynamic curriculum learning for imbalanced data classification,” in [*Proceedings of the IEEE/CVF International Conference on Computer Vision*], 5017–5026 (2019).
- [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] Zhang, 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] Zhang, D., Meng, D., Zhao, L., and Han, J., “Bridging saliency detection to weakly supervised object detection based on self-paced curriculum learning,” *arXiv preprint arXiv:1703.01290* (2017).
- [13] Zhang, Y., David, P., and Gong, B., “Curriculum domain adaptation for semantic segmentation of urban scenes,” in [*Proceedings of the IEEE international conference on computer vision*], 2020–2030 (2017).
- [14] 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).
- [15] Murali, A., Pinto, L., Gandhi, D., and Gupta, A., “Cassl: Curriculum accelerated self-supervised learning,” in [*2018 IEEE International Conference on Robotics and Automation (ICRA)*], 6453–6460, IEEE (2018).
- [16] Tang, Y., Wang, X., Harrison, A. P., Lu, L., Xiao, J., and Summers, R. M., “Attention-guided curriculum learning for weakly supervised classification and localization of thoracic diseases on chest radiographs,” in [*International Workshop on Machine Learning in Medical Imaging*], 249–258, Springer (2018).
- [17] Sarafianos, N., Giannakopoulos, T., Nikou, C., and Kakadiaris, I. A., “Curriculum learning for multi-task classification of visual attributes,” in [*Proceedings of the IEEE International Conference on Computer Vision Workshops*], 2608–2615 (2017).
- [18] Sarafianos, N., Giannakopoulos, T., Nikou, C., and Kakadiaris, I. A., “Curriculum learning of visual attribute clusters for multi-task classification,” *Pattern Recognition* **80**, 94–108 (2018).
- [19] Dong, Q., Gong, S., and Zhu, X., “Multi-task curriculum transfer deep learning of clothing attributes,” in [*2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*], 520–529, IEEE (2017).
- [20] Wang, C., Zhang, Q., Huang, C., Liu, W., and Wang, X., “Manacs: 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).

- [21] Gong, C., “Exploring commonality and individuality for multi-modal curriculum learning,” in [*Thirty-first AAAI conference on artificial intelligence*], (2017).
- [22] Jiang, L., Meng, D., Zhao, Q., Shan, S., and Hauptmann, A. G., “Self-paced curriculum learning,” in [*Twenty-Ninth AAAI Conference on Artificial Intelligence*], (2015).
- [23] Matiisen, T., Oliver, A., Cohen, T., and Schulman, J., “Teacher–student curriculum learning,” *IEEE transactions on neural networks and learning systems* **31**(9), 3732–3740 (2019).
- [24] Weinshall, D., Cohen, G., and Amir, D., “Curriculum learning by transfer learning: Theory and experiments with deep networks,” in [*International Conference on Machine Learning*], 5238–5246, PMLR (2018).
- [25] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al., “Pytorch: An imperative style, high-performance deep learning library,” *Advances in neural information processing systems* **32**, 8026–8037 (2019).
- [26] Simonyan, K. and Zisserman, A., “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556* (2014).