

LOOK AT HERE : UTILIZING SUPERVISION TO ATTEND SUBTLE KEY REGIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

Despite the success of deep learning in computer vision, algorithms to recognize subtle and small objects (or regions) is still challenging. For example, recognizing a baseball or a frisbee on a ground scene or a bone fracture in an X-ray image can easily result in overfitting, unless a huge amount of training data is available. To mitigate this problem, we need a way to force a model should identify subtle regions in limited training data. In this paper, we propose a simple but efficient supervised augmentation method called Cut&Remain. It achieved better performance in multi-class classification tasks (clavicle and pelvic X-ray) and a multi-label classification task of small objects (MS-COCO_s) than other supervised augmentation and the explicit guidance methods. In addition, using the class activation map, we identified that the Cut&Remain methods drive a model to focus on relevant subtle and small regions efficiently. We also show that the performance monotonically increased along the Cut&Remain ratio, indicating that a model can be improved even though only limited amount of Cut&Remain is applied for, so that it allows low supervising(annotation) cost for improvement.

1 INTRODUCTION

Deep learning and convolutional neural networks (CNNs) have been successfully applied in various fields, including medicine, visual inspection, and self-driving cars (Hemdan et al., 2020; Ayan & Ünver, 2019; Chen et al., 2020; 2019; Sa et al., 2017; Yahalomi et al., 2019; Lai, 2015; Wang et al., 2018; Tajbakhsh et al., 2020); however, the utilization of machine learning techniques to recognize a subtle region or a small object in a scene is still challenging (Kang et al., 2017; Dai et al., 2016; Herath et al., 2017; Wu et al., 2017; Zhou et al., 2019), especially when sufficient training data are not available on which we frequently encounter the difficulty. Despite the data insufficiency, a model is trained based only on cross-entropy loss to fit the model to the one-hot target label in most classification tasks. As a result, the model is easily overfitted to the training data. To mitigate this problem, regularization methods, including soft labels, temperature scaling, and standard augmentations (e.g., flip, rotation), can be used; however, these are not a kind of explicit methods, which allow the model to recognize subtle and small regions. Occasionally, such as radiological diagnosis and UAV detection(Stojnić et al., 2021), a model must be trained to look at a relevant region even by utilizing human-guided supervision for ensuring safety. We might consider this is similar to imitation learning(Hussein et al., 2017) in the field of reinforcement learning.

Figure 1 shows a representative X-ray image of a clavicle fracture. A discontinuity is a morphological property representing a clavicle fracture (red box). However, a lot of discontinuities can be observed in the image (yellow boxes). In contrast to machine learning algorithms, radiologists can correctly diagnose among these discontinuities based on their background knowledge. In the absence of background knowledge, implicit embedding the feature representation of a clavicle fracture is very challenging. It can be solved by considering existing region-perturbed augmentations (Krizhevsky et al., 2012; DeVries & Taylor, 2017; Zhang et al., 2017; Yun et al., 2019) or attention-guided methods(Yang et al., 2019; Li et al., 2018) integrated with human supervision.

In this paper, we propose a simple but efficient augmentation strategy, called Cut&Remain, to allow a model to recognize subtle and small key regions. Figure 2 shows the augmented data and its methodology. Relevant area remains and rest area is zero-out, so that the positional information of the lesion and the image size are preserved. Especially in radiological diagnosis, the positional

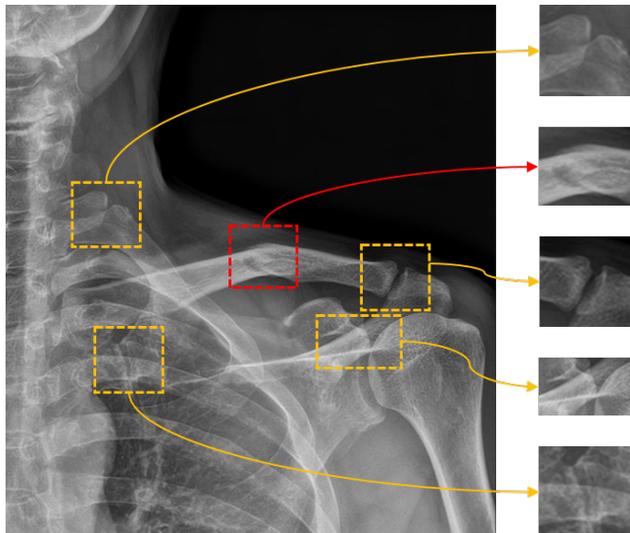


Figure 1: Example of candidate features of a lesion. The ground truth and the confusing background region are denoted as red and yellow bounding boxes, respectively.

information could be a cue for the nature of an area, such as intensity and scale of a lesion. Due to the general protocol of medical image acquisition, the anatomical structures are usually aligned so that the target lesion tends to be distributed at a specific position in the image. We show that taking this information into account can improve performance.

The performance of the Cut&Remain in binary (clavicle X-ray dataset), multiclass (pelvic X-ray dataset), and multilabel (MS-COCO_s, dataset composed of the small objects in MS-COCO 2017(Lin et al., 2014)) classification tasks was higher than that of other supervised methods, despite the limited training data and relatively small size of lesions or objects. In addition, the authors qualitatively analyzed the focus region of these models by using Grad-CAM (Selvaraju et al., 2017), and the results showed that the proposed model effectively focused on the relevant regions. Furthermore, the method was tested based on the number of annotations to investigate the effort required in this task. As a result, its performance monotonically improved with the number of annotations, indicating that, even when limited annotations are available, this method can help the model to learn relevant regions when compared to the baseline.

2 RELATED WORKS

2.1 REGION-PERTURBED AUGMENTATION TECHNIQUES

The cutout method (DeVries & Taylor, 2017), in which square regions of the input image are randomly masked out, has been proposed to improve the robustness of CNNs; however, information-rich pixels may be lost during training, which can be critical if the lesion size of the image is relatively small.

Meanwhile, the mixup method (Zhang et al., 2017) considers two samples, in which the ground truth label of the new sample is obtained by a combination of one-hot labels. These samples, however, confuse the model due to their ambiguity and unnaturalness.

The cutmix method (Yun et al., 2019), which is a crossover of the cutout and mixup methods, is a novel data augmentation strategy that compensates for the disadvantages of the cutout and mixup strategies. Cutmix produces new samples by cutting and pasting patches within minibatches, thereby enhancing the performance in many computer vision tasks. However, similarly to the cutout strategy, the loss of informative features of the lesion might result in performance degradation.

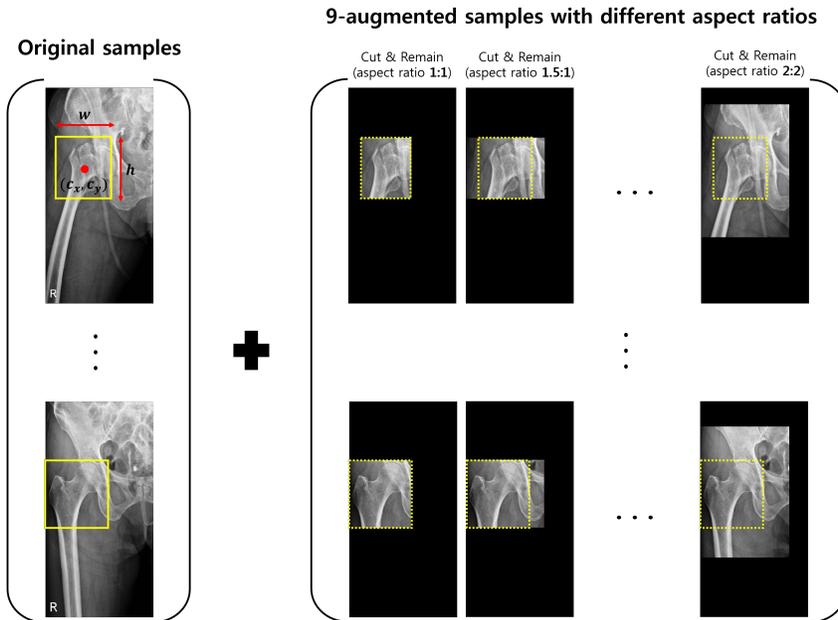


Figure 2: Mini-batch configuration using Cut&Remain data augmentation during training. The yellow box refer to bounding box annotation $\mathbf{B} = (c_x, c_y, w, h)$. The augmented samples with different aspect ratios are added to current mini-batch.

These methods are useful in combining multiple images or their cropped versions to create new training data; however, they are still not object-aware and have not been designed specifically for small object recognition. Copy-Paste (Ghiasi et al., 2021) is a simple strategy that combines information from multiple images, in an object-aware manner, to copy instances of objects from one image and paste them onto another. However, such an augmentation method for segmentation is difficult to apply in cases where the boundary of the region is anatomically ambiguous, such as fracture and anterior-superior iliac spine.

2.2 ATTENTION-GUIDED NETWORK

End-to-end trainable attention mechanisms that allow the network to focus on salient regions for image classification have been studied in the literature. Classification and attention mining streams have been previously used to explicitly model the attention mechanism during training (Li et al., 2018; Yang et al., 2019), and the attention map has been generated from the classification branch using gradient-based Grad-CAM (Selvaraju et al., 2017). These methods require an additional region guidance architecture, whereas Cut&Remain requires only the base network.

3 METHODS

Fundamentally, the proposed method utilizes human-annotated bounding box to zero-out unimportant areas in a image while preserving the spatial location of the important areas.

Assuming that $x \in \mathbb{R}^{W \times H \times C}$ and y denote a training image and its label, respectively, the goal of Cut&Remain is to generate a new training sample (\tilde{x}, \tilde{y}) , which can be described by Eq.(1) and used to train the model based on its original loss function.

$$\begin{aligned} \tilde{x} &= \mathbf{M} \odot x, \\ \tilde{y} &= y \end{aligned} \tag{1}$$

where $\mathbf{M} \in \{0, 1\}^{W \times H}$ denotes a binary mask indicating lesion and \odot is element-wise multiplication. To generate Mask \mathbf{M} , a bounding box annotation, $\mathbf{B} = (c_x, c_y, w, h)$ indicating the region to

	Augmented image \tilde{x}	Augmented label \tilde{y}
Sup-Mixup	$\lambda \mathbf{M} \odot x_A + (1 - \lambda) \mathbf{M} \odot x_B$	$\lambda y_A + (1 - \lambda) y_B$
Sup-Cutmix	$\mathbf{M} \odot x_A + (\mathbf{1} - \mathbf{M}) \odot x_B$	y_A
Sup-Cutout	Remove random region in $(\mathbf{1} - \mathbf{M})$	y

Table 1: Supervised-version of augmentation operations. We denote $\mathbf{1}$ as a binary mask and λ as the combination ratio. \mathbf{M} is the binary mask whose value is 1 inside the annotated box \mathbf{B} .

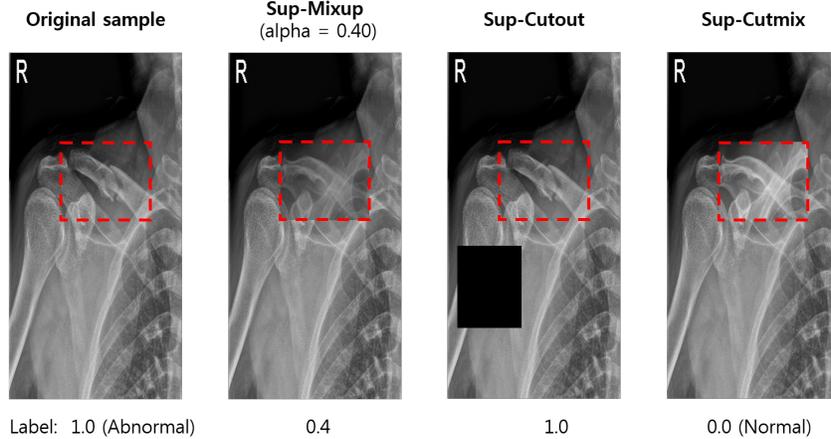


Figure 3: Overview of the augmented images from the baseline (Cutout, Mixup, and Cutmix) based on a supervised procedure on the Clavicle X-ray dataset. The red boxes refer to bounding box annotations.

remain in image x was used, where c_x, c_y, w, h indicate a box’s center coordinate, width and height respectively. To get more augmentations, the aspect ratio of $\{1.0, 1.5, 2.0\}$ is applied to box’s width and height so that results in nine boxes per single annotation. Then, the element of binary mask $\mathbf{M} \in \{0, 1\}^{W \times H}$ is set to 1 or 0 according to if it is inside or outside of \mathbf{B} . In this way, Cut&Remain keeps positional information and scale unchanged for relevant regions, that is distinguished from “cropping” that conducts resizing after cutting an area.

In each training step, an augmented sample (\tilde{x}, \tilde{y}) is generated based on each training sample according to Eq. (1) with different aspect ratios, and included together with original samples in the mini-batch, as shown in Figure 2.

4 EXPERIMENTS

We evaluated the Cut&Remain method considering internally sourced X-ray datasets for clavicle abnormality and femur fracture classification and the COCO 2017 dataset (Lin et al., 2014) for multi-label classification. A digital radiographic examination (CKY Digital Diagnost; Philips, Eindhoven, The Netherlands) included anteroposterior views obtained from patients. For the internal datasets, experienced surgeons annotated the bounding boxes of the relevant areas, including normal images.

Conventional augmentations were evaluated based on a supervised procedure. For each training image, we generated a binary mask \mathbf{M} . In each training step, the supervised version of the mixup (sup-mixup) and cutmix (sup-cutmix) augmented sample (\tilde{x}, \tilde{y}) was generated by mixing or combining two randomly selected training samples, (x_A, y_A) and (x_B, y_B) , in a mini-batch. A supervised cutout (sup-cutout) augmented sample was generated by removing a random region in $(\mathbf{1} - \mathbf{M})$, where $\mathbf{1}$ corresponds to a binary mask. The augmentation operations are listed in Table 1, and examples of each augmented image are shown in Figure 3.

The experimental results show that Cut&Remain outperforms other data augmentation and attention-guided network (AGN) techniques for all experiments. The AGN to which this method was com-

Method	AUC-ROC	F1-score
ResNet-50	90.8±1.5	91.2±1.8
Sup-Mixup	89.0±1.7	90.4±2.0
Sup-Cutout	95.6±1.2	95.0±1.4
Sup-Cutmix	96.8±1.7	97.2±1.6
AGN(Li et al., 2018)	93.4±1.1	93.0±1.4
Cut&Remain (w/o aspect ratio variation)	98.4±0.4	98.6±0.5
Cut&Remain (w/ aspect ratio variation)	98.6±0.4	98.8±0.4

Table 2: Results of the clavicle X-ray dataset. The averaged AUC-ROC and F1-score and their standard deviation of the 5-fold cross-validation are reported for each set.

pared had the same architecture as those proposed in Li et al. (2018), including additional supervision, while eliminating the attention mining branch. For the COCO 2017 dataset, several categories composed of small objects were selected.

4.1 BINARY CLASSIFICATION ON THE CLAVICLE X-RAY DATASET

We collected 1,080 clavicle X-rays, including 322 abnormal images (270 fractures and 52 dislocations). The dataset, which was collected from a real-world cohort, also contained cases in which other abnormalities (not only clavicle fracture or dislocation), were present. The overall images were vertically cut in half to increase the number of samples and reduce the resolution of the images. The dataset was split into training (80%), validation (10%), and test (10%) sets. The area under the receiver operating characteristics (AUC-ROC) (Fawcett, 2006) and the F1-score (Sasaki et al., 2007), obtained after 5-fold cross validation to investigate the classification performance, is reported in this study. We selected ResNet-50 (He et al., 2016) as the backbone network, and its weights were randomly initialized. Binary cross-entropy loss was used in the classification and the momentum of the employed SGD optimizer was 0.9. The initial learning rate was set to 0.01. The model was trained until 2,000 epochs were performed, and the learning rate was reduced by a factor of 10 at epoch 1,500.

4.1.1 RESULTS

The experimental results are listed in Table 2. Compared to the supervised version of other augmentation techniques, Cut&Remain showed the highest AUC-ROC (98.6) and F1-score (98.8), and the performance of the base network was improved by more than 7.8 and 7.6, respectively. The results demonstrate that the training of CNNs using random region mixing, removing, or replacing may be affected by irrelevant noisy areas. As a result, bias towards a noisy distractor might be added, harming the generalization performance when the same correlations of the test data are not present in the training data. In contrast, the subtle differences can be exploited in local structures when employing Cut&Remain by highlighting the object of interest and suppressing irrelevant areas. Meanwhile, the utilization of sup-mixup led to the degradation of the performance of the algorithms in our dataset, as it might provide unnatural artifacts caused by the overlap between tissues. The results of the Cut&Remain support the hypothesis that the lesion must be considered in images that depict a small lesion.

4.2 MULTICLASS CLASSIFICATION OF FEMUR FRACTURE ON THE PELVIC X-RAY DATASET

A total of 740 anteroposterior pelvic X-ray images were acquired from our institution. Two experienced surgeons reviewed the cases and identified 380 fracture cases and 360 normal cases, following the Arbeitsgemeinschaft Osteosynthese foundation/Orthopedic Trauma Association (AO/OTA) classification standard, in the overall dataset. The dataset was divided into 360 normal (nonfracture), 273 A-type (trochanteric region), and 107 B-type (neck) cases based on the presence of a fracture and its position. The overall images were vertically cut in half to increase the number of samples and reduce the resolution of the images. We split the dataset into 924, 102, and 454 images for training, validation, and testing, respectively, and recorded the AUC-ROC and F1-score for evaluating the classification performance. In all experiments, the ResNet-50 (He et al., 2016) was selected as

Method	AUC-ROC			F1-score		
	Normal	A-type	B-type	Normal	A-type	B-type
ResNet-50	91.2±0.8	91.2±1.3	87.2±1.8	92.2±0.9	73.0±1.1	47.4±2.5
Lee et al. (2020)	90.6±1.8	91.0±1.0	88.4±1.5	95.4±1.4	88.2±1.8	76.4±2.9
Sup-Mixup	90.0±1.5	85.8±0.9	86.4±0.7	92.0±1.7	67.2±1.3	63.8±3.9
Sup-Cutout	97.0±1.3	96.4±1.5	93.0±1.8	96.4±1.1	85.0±1.7	68.6±3.6
Sup-Cutmix	95.8±1.6	94.8±1.3	93.0±1.7	96.0±2.0	83.4±2.2	77.0±3.6
AGN (Li et al., 2018)	94.6±0.6	92.8±0.8	90.2±0.8	94.8±1.3	84.4±3.5	66.2±3.4
Cut&Remain (w/o aspect ratio variation)	97.4±1.3	97.0±1.4	97.0±1.1	98.4±1.2	93.0±1.4	85.8±2.7
Cut&Remain (w/ aspect ratio variation)	97.8±0.8	97.4±1.1	97.2±1.3	98.6±0.8	93.8±1.2	87.0±2.3

Table 3: Results of the pelvic X-ray dataset. The averaged AUC-ROC and F1-score and their standard deviation of the 5-fold cross-validation are reported for each set.

our backbone network, and its weights were randomly initialized. We used cross-entropy loss for classification and set the momentum of the employed SGD optimizer to 0.9. The initial learning rate was set to 0.001. The model was trained for 2,000 epochs, and the learning rate dropped by a factor of 10 after epoch 1,500.

4.2.1 RESULTS

The experimental results are listed in Table 3. On pelvic X-ray images, the AUC-ROC and F1-score obtained after applying Cut&Remain to ResNet-50 were improved by 6.6, and 6.4 for Normal-class classification, respectively. Lee et al. (2020) presented a method for classifying femur fractures on X-ray images using deep learning trained with radiology reports. In the literature, they achieved an average F1 score of 81.7 in the 3-class classification task when using the whole images, which is not favorable to translate clinical practice. Conversely, the results of the present study indicate that Cut&Remain can significantly improve the classification performance due to the high F1 scores and AUCs achieved using this method.

4.3 MULTI-LABEL CLASSIFICATION ON THE COCO 2017 DATASET

In this experiment, for generality, we validated the performance of the Cut&Remain in the natural image domain. The Microsoft COCO 2017 (Lin et al., 2014) is a standard dataset built by Microsoft for object detection, image segmentation, and other applications. The training set was composed of 118,287 images that depicted common objects in scenes. These objects are categorized based on 80 classes, and every image contains an average of 2.9 objects. As the ground-truth label of the test set is not available, we evaluated a validation set of 5,000 images. The dataset includes various types of small objects and complex backgrounds; therefore, it is suitable for small object classification. Only images in which the average area occupied by the object is lower than 2% of the whole image size were added to the dataset to obtain only the images containing small objects. Based on this criteria, we selected 66,612 images for training and 2,805 images for testing that included 27 types of objects, including cars, traffic lights, kites, and cups, to create a training dataset, which was named MS-COCO_s (see the supplement for details). The mean average precision (mAP), average per-class F1 (CF1), and average overall F1 (OF1), which are the same statistical variables reported in the conventional settings for COCO 2017 (Liu et al., 2018; Wang et al., 2020), were calculated. For all experiments, we employed ResNet-50 (He et al., 2016) as our backbone network, whose weights were randomly initialized. We used binary cross-entropy loss for classification and an Adam optimizer. The initial learning rate was set to 0.001. The model was trained for 50 epochs, and the learning rate dropped by a factor of 10 at epoch 40.

4.3.1 RESULTS

Table 4 indicates that Cut&Remain provided better results than those of the baseline on MS-COCO_s, whereas the mAP was 2.3% higher than that of the baseline network. However, the performance

Method	mAP	CF1	OF1
ResNet-50	74.4	69.3	74.0
Sup-Mixup	73.5	68.6	73.0
Sup-Cutout	75.2	69.1	74.0
Sup-Cutmix	75.6	71.1	75.0
AGN(Li et al., 2018)	76.2	71.0	75.7
Cut&Remain (w/o aspect ratio variation)	75.8	70.8	75.1
Cut&Remain (w/ aspect ratio variation)	76.7	71.8	75.6

Table 4: Multi-label classification results on the MS-COCO_s testset. All metrics are in %. Results are reported for an input resolution of 448.

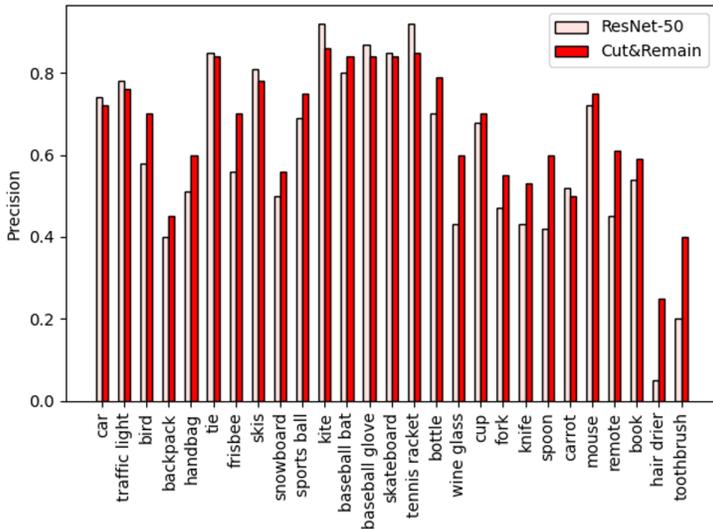


Figure 4: The class-wise precision on the MS-COCO_s dataset. The bars represent the results achieved by using baseline and Cut&Remain, respectively. Cut&Remain improves results on difficult classes, like hair drier, backpack, and toothbrush

improvement using Cut&Remain was relatively small compared to the previous tasks, which considered medical images, because objects are distributed in various locations in the natural image domain. As a result, using the location information of the lesion is difficult.

Although originally developed for medical image tasks, the classification performance in the natural image domain was also improved after applying Cut&Remain. The results imply that the object context is crucial for X-ray images and generally helpful for distinguishing subtle differences in local structures in the natural image domain. The class-wise precision on the MS-COCO_s is shown in Figure 4.

4.4 WHAT DOES A MODEL LEARN WITH CUT&REMAIN?

4.4.1 QUALITATIVE EVALUATION BY GRAD-CAM

We have validated Cut&Remain such that lesions are mainly considered as cues for classification and the motivation shared by attention-guided networks. To verify that Cut&Remain recognizing the key lesion on target images after the learning procedure, the activation maps of the test images trained by Cut&Remain were compared to those trained with the Sup-Cutout, Sup-Mixup, and Sup-Cutmix. Figure 5 shows the test examples and the corresponding Grad-CAM for the abnormal classes. We used the vanilla ResNet-50 model to obtain the Grad-CAM and observe the effect of only using the

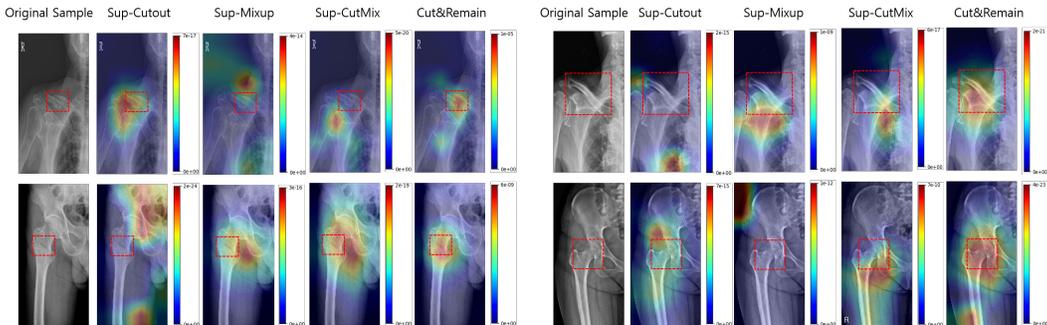


Figure 5: Grad-CAM visualization on test images using the model trained with each augmentation technique. Left: Grad-CAM for abnormal class. Right: Grad-CAM for normal class. Ground-truth annotation are shown as a red boxes.

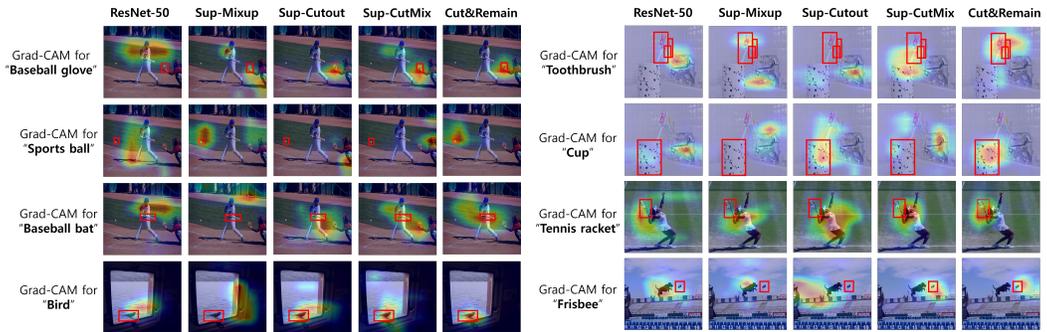


Figure 6: Grad-CAM visualization on MS-COCO_s test images. Ground-truth annotation are shown as a red boxes.

augmentation method. For the images in which no abnormalities were present, Cut&Remain could also learn the corresponding mask, thus distinguishing normal from other abnormal classes.

We observed that Cut&Remain allowed a model to detect the lesions. In contrast, mixing introduces unnatural artifacts; therefore, the corresponding model was confused when choosing cues for recognition, as shown in the Grad-CAM, which might lead to suboptimal classification performance, as presented in Tables 2 and 3.

We also visualize the attended regions in the MS-COCO_s dataset, as shown in Figure 6. Cut&Remain could also learn the corresponding context in the natural image domain. Examples containing more samples are presented in the supplementary material.

4.4.2 FEATURE REPRESENTATION PROPERTY

To verify that Cut&Remain assists in generating a background-independent feature vector (i.e., whether the model mainly focuses on the lesion), we analyzed the similarity of the vectors produced by using the original and augmented images. We compared the ResNet-50 trained without any augmentation techniques and individually trained using the Sup-Cutout, Sup-Mixup, Sup-Cutmix, and Cut&Remain strategies. We conducted this investigation based on the training set of clavicle X-rays and the same experimental setting as the subsection of “Binary classification on the clavicle X-ray dataset”. The Euclidean distance and cosine distance were the similarity measures calculated considering the feature vectors obtained from the output of the last convolution layer.

The experimental results are presented in Table 5. We observed that Cut&Remain created similar representation vectors for the original and augmented samples. In contrast, conventional augmentation techniques, which randomly remove, mix, or replace regions in images, resulted in increased dissimilarity due to the informative features that might have been lost.

	Sup-Cutout	Sup-Mixup	Sup-Cutmix	Cut&Remain
Euclidean distance	35.60±16.34	40.56±25.12	37.53±17.82	27.45±13.31
Cosine distance	0.73±0.16	0.77±0.18	0.79±0.17	0.63±0.16

Table 5: Average feature vector similarity between the original and augmented samples by a supervised version of the cutout, mixup, cutmix, and Cut&Remain. Each vector was obtained from the last convolution layer of the ResNet-50 trained with the clavicle dataset

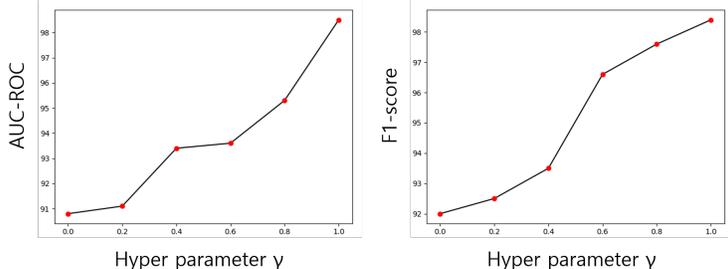


Figure 7: Effect of the Cut&Remain and augmented images in the training dataset (in Clavicle X-ray dataset)

4.4.3 PERFORMANCE ACCORDING THE NUMBER OF ANNOTATIONS

We evaluated the performance according amounts of the training dataset applying Cut&Remain. The ratio γ of the training data applying Cut&Remain was $\{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$. The performance of Cut&Remain with different γ is given in Figure 6. The Cut&Remain achieved the best performance after augmentation was applied to the overall training dataset (i.e., $\gamma = 1.0$). Furthermore, the AUC-ROC and F1-score monotonically increased with γ ; therefore, the performance improvement is guaranteed even if limited number of annotations are available.

5 CONCLUSION

We introduced a novel data augmentation strategy, namely Cut&Remain, for training CNNs. This strategy has a strong attention-guided benefit, can be easily implemented, and has no computational overhead. Furthermore, it is surprisingly effective on medical image datasets. We have shown that Cut&Remain can significantly improve the performance of the CNN classifier on various image domains, despite the limited amount of training data and relatively small lesion size. In particular, on clavicle X-ray classification, the AUC-ROC and F1-score of the ResNet-50 were improved by 7.8 and 7.6, respectively, when the Cut&Remain method was applied. On Femur fracture classification, Cut&Remain resulted in an AUC-ROC improvement of 6.6 for Normal-class classification when compared to that of the baseline. The explicit attention mechanism, however, did not guarantee performance improvement. In the natural image domain, MS-COCO_s, this data augmentation technique provided consistent improvements over the baseline and other augmentation techniques. Moreover, Grad-CAM analysis and t-SNE visualization indicated that Cut&Remain resulted in more focus on lesions, irrespective of the background.

REFERENCES

- Enes Ayan and Halil Murat Ünver. Diagnosis of pneumonia from chest x-ray images using deep learning. In *2019 Scientific Meeting on Electrical-Electronics & Biomedical Engineering and Computer Science (EBBT)*, pp. 1–5. Ieee, 2019.
- Hu Chen, Kailai Zhang, Peijun Lyu, Hong Li, Ludan Zhang, Ji Wu, and Chin-Hui Lee. A deep learning approach to automatic teeth detection and numbering based on object detection in dental periapical films. *Scientific reports*, 9(1):1–11, 2019.
- Kai-Chi Chen, Hong-Ren Yu, Wei-Shiang Chen, Wei-Che Lin, Yi-Chen Lee, Hung-Hsun Chen, Jyun-Hong Jiang, Ting-Yi Su, Chang-Ku Tsai, Ti-An Tsai, et al. Diagnosis of common pulmonary diseases in children by x-ray images and deep learning. *Scientific reports*, 10(1):1–9, 2020.
- Jifeng Dai, Kaiming He, and Jian Sun. Instance-aware semantic segmentation via multi-task network cascades. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3150–3158, 2016.
- Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- Tom Fawcett. An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874, 2006.
- Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2918–2928, 2021.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Ezz El-Din Hemdan, Marwa A Shouman, and Mohamed Esmail Karar. Covidx-net: A framework of deep learning classifiers to diagnose covid-19 in x-ray images. *arXiv preprint arXiv:2003.11055*, 2020.
- Samitha Herath, Mehrtash Harandi, and Fatih Porikli. Going deeper into action recognition: A survey. *Image and vision computing*, 60:4–21, 2017.
- Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017.
- Kai Kang, Hongsheng Li, Junjie Yan, Xingyu Zeng, Bin Yang, Tong Xiao, Cong Zhang, Zhe Wang, Ruohui Wang, Xiaogang Wang, et al. T-cnn: Tubelets with convolutional neural networks for object detection from videos. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(10):2896–2907, 2017.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- Matthew Lai. Deep learning for medical image segmentation. *arXiv preprint arXiv:1505.02000*, 2015.
- Changhwan Lee, Jongseong Jang, Seunghun Lee, Young Soo Kim, Hang Joon Jo, and Yeesuk Kim. Classification of femur fracture in pelvic x-ray images using meta-learned deep neural network. *Scientific reports*, 10(1):1–12, 2020.
- Kunpeng Li, Ziyang Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. Tell me where to look: Guided attention inference network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9215–9223, 2018.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.

- Yongcheng Liu, Lu Sheng, Jing Shao, Junjie Yan, Shiming Xiang, and Chunhong Pan. Multi-label image classification via knowledge distillation from weakly-supervised detection. In *Proceedings of the 26th ACM international conference on Multimedia*, pp. 700–708, 2018.
- Ruhan Sa, William Owens, Raymond Wiegand, Mark Studin, Donald Capoferri, Kenneth Barooha, Alexander Greaux, Robert Rattray, Adam Hutton, John Cintineo, et al. Intervertebral disc detection in x-ray images using faster r-cnn. In *2017 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 564–567. IEEE, 2017.
- Yutaka Sasaki et al. The truth of the f-measure. 2007, 2007.
- 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*, pp. 618–626, 2017.
- Vladan Stojnić, Vladimir Risojević, Mario Muštra, Vedran Jovanović, Janja Filipi, Nikola Kezić, and Zdenka Babić. A method for detection of small moving objects in uav videos. *Remote Sensing*, 13(4), 2021. doi: 10.3390/rs13040653. URL <https://www.mdpi.com/2072-4292/13/4/653>.
- Nima Tajbakhsh, Laura Jeyaseelan, Qian Li, Jeffrey N Chiang, Zhihao Wu, and Xiaowei Ding. Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation. *Medical Image Analysis*, 63:101693, 2020.
- Guotai Wang, Wenqi Li, Maria A Zuluaga, Rosalind Pratt, Premal A Patel, Michael Aertsen, Tom Doel, Anna L David, Jan Deprest, Sébastien Ourselin, et al. Interactive medical image segmentation using deep learning with image-specific fine tuning. *IEEE transactions on medical imaging*, 37(7):1562–1573, 2018.
- Ya Wang, Dongliang He, Fu Li, Xiang Long, Zhichao Zhou, Jinwen Ma, and Shilei Wen. Multi-label classification with label graph superimposing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 12265–12272, 2020.
- Qi Wu, Chunhua Shen, Peng Wang, Anthony Dick, and Anton Van Den Hengel. Image captioning and visual question answering based on attributes and external knowledge. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1367–1381, 2017.
- Erez Yahalomi, Michael Chernofsky, and Michael Werman. Detection of distal radius fractures trained by a small set of x-ray images and faster r-cnn. In *Intelligent Computing-Proceedings of the Computing Conference*, pp. 971–981. Springer, 2019.
- Heechan Yang, Ji-Ye Kim, Hyongsuk Kim, and Shyam P Adhikari. Guided soft attention network for classification of breast cancer histopathology images. *IEEE transactions on medical imaging*, 39(5):1306–1315, 2019.
- Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 6023–6032, 2019.
- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127(3):302–321, 2019.