

Window Loss for Bone Fracture Detection and Localization in X-ray Images with Point-based Annotation

* † Xinyu Zhang¹, * Yirui Wang¹, Chi-Tung Cheng², Le Lu¹, Adam P. Harrison¹, Jing Xiao³, Chien-Hung Liao², Shun Miao¹

¹ PAII Inc., Bethesda, Maryland, USA

² Chang Gung Memorial Hospital, Linkou, Taiwan, ROC

³ Ping An Technology, Shenzhen, China

Abstract

Object detection methods are widely adopted for computer-aided diagnosis using medical images. Anomalous findings are usually treated as objects that are described by bounding boxes. Yet, many pathological findings, *e.g.*, bone fractures, cannot be clearly defined by bounding boxes, owing to considerable instance, shape and boundary ambiguities. This makes bounding box annotations, and their associated losses, highly ill-suited. In this work, we propose a new bone fracture detection method for X-ray images, based on a labor effective and flexible annotation scheme suitable for abnormal findings with no clear object-level spatial extents or boundaries. Our method employs a simple, intuitive, and informative point-based annotation protocol to mark localized pathology information. To address the uncertainty in the fracture scales annotated via point(s), we convert the annotations into pixel-wise supervision that uses lower and upper bounds with positive, negative, and uncertain regions. A novel *Window Loss* is subsequently proposed to only penalize the predictions outside of the uncertain regions. Our method has been extensively evaluated on 4410 pelvic X-ray images of unique patients. Experiments demonstrate that our method outperforms previous state-of-the-art image classification and object detection baselines by healthy margins, with an AUROC of 0.983 and FROC score of 89.6%.

Introduction

Due to its efficiency, accessibility, and low cost, X-ray imaging is one of the most commonly performed diagnostic examinations in clinical practice. Conventional computer vision techniques have been extensively researched in the past few decades to develop computer aided diagnosis (CAD) solutions, aiming to increase image reading efficiency and reduce misdiagnoses risks. Current state-of-the-art CAD systems widely adopt generic principles of image classification and object detection to identify and localize the anomalies. However, to provide both robust classification and localization in applications where pathologies have ambiguously defined extents, such as for pelvic X-rays (PXR), alternative solutions are needed. This is the topic of our work.

* X. Zhang and Y. Wang—Equal contribution.

† This work was done when X. Zhang interned at PAII Inc.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

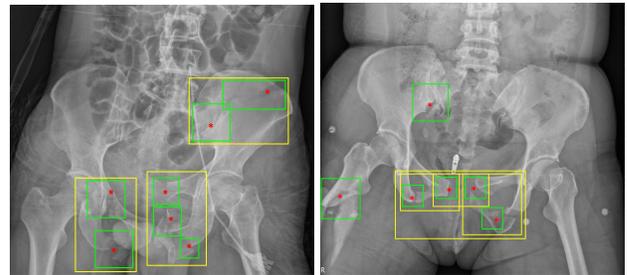


Figure 1: Fractures in PXR where the numbers of instances and their spatial extents are extremely ambiguous. The annotations of fractures can be represented by small/tight bounding boxes on fracture lines (green boxes), or relatively big bounding boxes containing the entire fractured region (yellow boxes). To skirt this inherent subjectivity, we adopt and learn from point-based annotations (red dots).

Spurred by the public release of large-scale datasets, chest X-ray (CXR) CAD has received a great deal of recent attention, with many prominent methods treating the problem as a multi-label image classification problem (Rajpurkar et al. 2017; Wang et al. 2017; Baltruschat et al. 2019). These works typically report area under receiver operating characteristic curve (AUROC) scores, *e.g.*, 0.817 (Baltruschat et al. 2019) on the NIH-ChestXray14 dataset (Wang et al. 2017) and 0.930 (Pham et al. 2020) on the CheXPert dataset (Irvin et al. 2019). Yet, in clinical diagnostic workflows, localization of the detected anomalies is also needed both to interpret the results and for verifiable decision making. While class activation maps (CAMs) can provide a degree of localization under classification setups, they have limited localization power and tend to miss abnormal sites when multiple pathological findings occur concurrently in an image.

To improve localization performance, several methods use additional box-level annotations to facilitate the training of image classifiers (Li et al. 2018; Liu et al. 2019a; Huang et al. 2020). For other CAD applications where localization performance is more central, such as lung nodule and breast mass/calcification detection, explicit object detection frameworks are widely adopted (Yan et al. 2018b, 2019; Jiang et al. 2020). However, compared to image classifiers, object detection methods focus predominantly on localization and

may struggle to produce robust image-level diagnoses, *e.g.*, by triggering undesirable false positives.

Another critical downside to object detectors is their reliance on bounding boxes. Besides being prohibitively laborious, bounding boxes (or any explicit defining boundary) are not always suitable for all types of medical anomalies, particularly when pathological findings cannot be easily delineated as discrete instances with definite boundaries. Fractures in PXR serve as excellent exemplar, as they usually involve multiple anatomical sites with complex morphology and curvilinear fragments that make defining the border, extent, and number of the bounding boxes extraordinarily difficult and ambiguous (as shown in Fig. 1). Therefore, an anomaly diagnosis and localization method with 1) a cost-effective and flexible annotation scheme and 2) simultaneously high classification and localization performance is critical and highly desirable.

In this work, we present a trauma PXR bone fracture diagnosis and localization system with a cost-effective and flexible annotation scheme that is suitable for detecting abnormal findings with no clear object boundaries. Interpreting trauma PXRs in an emergency room environment requires timely image-level diagnoses and fracture localizations. The image-level diagnosis determines the downstream actions in the clinical workflow, *e.g.*, further examination, hospitalization, or hospital discharge. Robust fracture localization can help physicians avoid missed diagnoses that must be treated in a timely manner, otherwise serious complications can occur. Due to the geometrical complexities of bone fractures and the perspective projection distortions of PXRs, diagnosis and localization of fractures in PXR is a challenging task where even experienced physicians can make considerable diagnostic errors (Chellam 2016).

To cope with the inherent geometrical complexity of pelvic fractures, our model use annotated points as supervisory signals. This point-based annotation is labor efficient and allows ambiguous fracture conditions to be adequately represented as one point or multiple points (on the fracture sites). To account for the uncertainties in the scale and shape of the annotated fracture sites, we propose to convert the points to pixel-wise supervision signals, which consist of lower and upper bounds of the expected network prediction with an allowed region between the bounds. A novel *Window Loss* is proposed that only penalizes predictions outside of the bounds, which encourages the model to learn from the reliable information without being distracted by the inherent uncertainties in the annotations. We evaluated our proposed method on 4410 PXRs collected from a primary trauma center representing real clinical situations. Our method reports an AUROC of 0.983 and FROC score of 89.6%, significantly outperforming previous state-of-the-art image classification and object detection methods by at roughly 1% or more on AUROC and 1.5% on FROC score, respectively.

Related Work

Image Classification Methods

The availability of large-scale public CXR datasets with image-level labels has catalyzed a large body of work

on CXR classification. Typically, these methods train deep neural networks that output per-image classification scores and calculate losses against the image-level labels. *CheXNet* (Rajpurkar et al. 2017) trains a DenseNet-121 (Huang et al. 2017) with global average pooling (GAP) performed on the penultimate feature maps before being inputted into a fully connected layer, and Wang et al. (2017) replaced the GAP with log-sum-exp (LSE). Yao et al. (2018); Wang et al. (2019); Li et al. (2018); Liu et al. (2019a) take a different approach by regarding the convolutional feature maps as local appearance descriptors of spatially sliced blocks. The features maps are then processed with a 1x1 convolution to produce class-wise probability maps and globally pooled to produce the image-level classification probabilities. Choices of pooling methods include LSE pooling (Yao et al. 2018; Wang et al. 2019) or the product of the probabilities (Li et al. 2018; Liu et al. 2019a), where the latter assumes that probability values are independent and identically distributed.

Methods with global feature pooling (Yan et al. 2018a) typically rely on spatial attention mechanisms like CAMs or Grad-CAM to localize the underlying anomalies. In contrast, methods with class-wise probability maps (Yao et al. 2018; Wang et al. 2019) can directly produce localizations without any add-on components. Either way, deep models trained with only image-level labels may not always produce plausible localizations with high reliability. Other approaches (Li et al. 2018; Liu et al. 2019a) have addressed this issue by developing strategies to exploit the small number/percentage of images with bounding box annotations.

Similar to prior work (Chen et al. 2020), our method uses a 1x1 convolution on the last feature map to produce a fracture class probability map. Different from the above classification based methods, our method is trained using point-based labels. In this way, we share a similar motivation with Li et al. (2018); Liu et al. (2019a) to leverage additional low cost localization annotations. However, we apply a fundamentally different conceptual framework, using 1) point-based annotations to account for the intrinsic ambiguity on visually defining/depicting bone fractures in X-rays as spatially-bounded objects and 2) designing a new Window Loss to naturally handle and learn from these annotations.

Object Detection Methods

Substantial progress has been made in recent years in the development of deep learning based object detection methods. These include anchor-based solutions (Ren et al. 2015; Lin et al. 2017b) and recent anchor-free solutions (Tian et al. 2019; Law and Deng 2018), which repeatedly break the accuracy records on benchmarks like MS-COCO (Lin et al. 2014a). The success and popularity of object detection in natural imagery has motivated research on adopting them for a wide spectrum of medical image anomaly localization problems, such as distal radius fractures in hand X-rays (Yahalomi, Chernofsky, and Werman 2019), pneumonia in CXR (Sirazitdinov et al. 2019), masses in mammography (Al-Masni et al. 2018) and general lesions in computed tomography (CT) (Yan et al. 2018b, 2019). Many domain specific enhancements have been proposed, including incor-

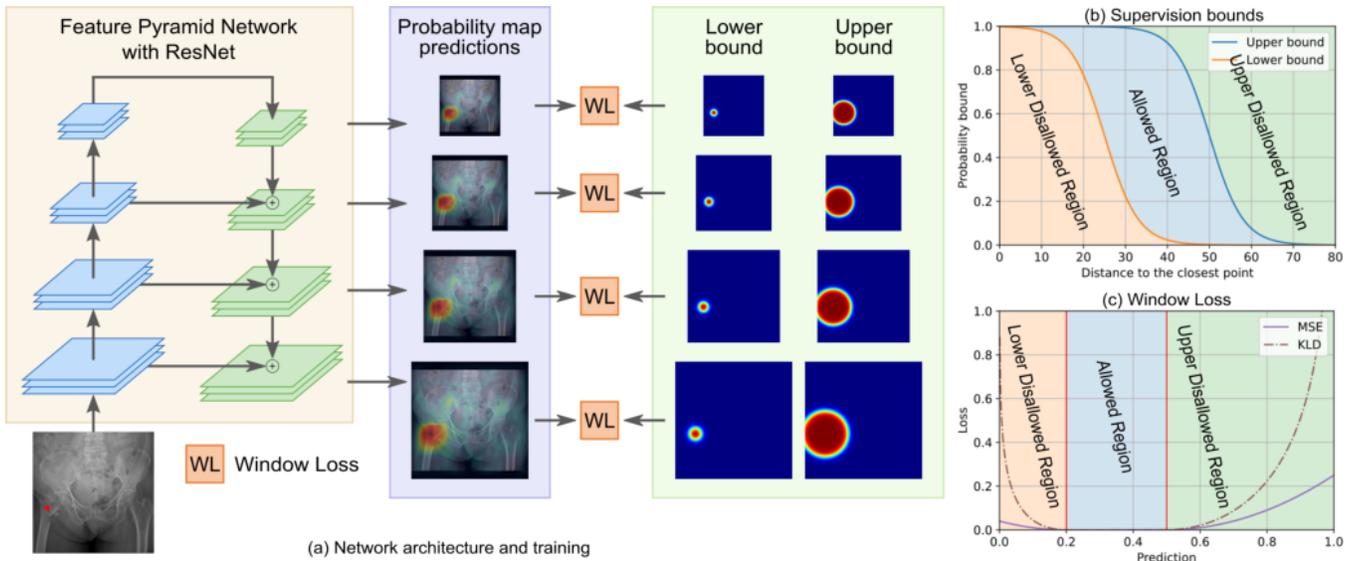


Figure 2: Framework of the proposed method. (a) Network architecture and training mechanism of the proposed method. (b) Lower and upper bounds of the generated pixel-level supervision. (c) Profiles of the proposed Window Loss using MSE or KLD as the divergence measure.

porating symmetry cues for breast mammography (Liu et al. 2019b); exploiting the anatomical prior of vertebrae via self-calibration (Zhao et al. 2019); and addressing missing annotations, hard negatives, and/or heterogeneous datasets for universal lesion detection (Yan et al. 2020; Cai et al. 2020). Like object detectors, our work explicitly learns from localization annotations, but the inherent visual ambiguity of pelvic fractures makes our problem fundamentally different. To address this, we propose a point-based annotation scheme and Window Loss for such challenging ambiguities.

Method

There are two critical metrics to measure a CAD solution’s effectiveness in clinical settings: **generality** (handling of unseen patients) and **precision** (generating good accuracy robustly for the given task). Naturally, incorporating as many patients as possible for training is desirable, but patient data often cannot be fully annotated at scale, precluding approaches seen in natural images (Lin et al. 2014b). In medical imaging, the labor costs are always paramount. Even more intractably, for many applications, *e.g.*, PXR fracture detection, there are inherent perceptual uncertainties and challenges in annotating precise bounding boxes of pathological findings. Alternative strategies are needed. A promising inspiration can be found in extreme points for object annotation (Papadopoulos et al. 2017), which are found to be a quicker and easier visual perception protocol by human annotators. Similarly, in this work, we use only point-based annotations to localize bone fractures in PXRs (*i.e.*, *where to look?*) so we can execute the annotation process at scale to cover as many as thousands of patients to achieve **generality**. To achieve **precision**, we propose a Window Loss that can robustly cope with the aforementioned ambiguities in defining the extents and scales of fracture sites (*i.e.*, *how*

to look?). Fig. 2 illustrates an overview of our framework.

Point-based Annotation

The complex morphology of pelvic fractures challenges the collection of reliable and consistent localization annotations. Due to the difficulties in defining the border, extent, and number of bounding boxes, the commonly adopted bounding box annotation suffers from high inter-annotator variations, which can cause conflicting supervisory signals. For instance, the pelvic fractures shown in Fig. 1 can be annotated using multiple small boxes or one larger box. We propose to use point-based annotations to provide the localization information. The pin-pointing annotation protocol is flexible and naturally fits our complex scenario: we ask the annotators/physicians to place points on visible fracture sites. For complex scenarios where the instance of fractures cannot be clearly defined, the annotators decide to place one or multiple points at their own discretion. Two challenges arise when using point-based annotation to supervise network training: 1) unknown scales of fracture sites, and 2) inter-rater variability in point placements. We address this with our proposed Window Loss.

Window Loss

Despite any uncertainties in scale and inter-annotator variability, point-based annotations provide the following information: 1) regions in close proximity to the annotated points are likely abnormal; 2) regions far away from all points are highly likely to be normal; 3) for in-between regions it is difficult, if not possible, to specify any particular confidence value. We encode this by calculating an allowed confidence range at every point in the image. Any prediction that falls into the allowed range is deemed acceptable, regardless of where it falls. This aligns with the inherent am-

biguities of fracture localization, where many image regions do not permit the specification of any particular value. On the other hand, if any prediction violates the allowed confidence range, it is penalized.

To calculate an allowed confidence range, we create two heatmaps to represent the upper confidence and lower confidence bounds at every pixel location. Both heatmaps are set to 1 at any point annotation and their values decline toward 0 at points further away. However, the lower confidence bound reduces rapidly as the distance to the point increases, whereas the upper bound reduces more gradually. We construct both heatmaps using the sigmoid of the distance to the annotation points:

$$\ell(i, j) = \max_k \sigma\left(\frac{r_\ell - D_{ij}^k}{\tau}\right), \quad (1)$$

$$u(i, j) = \max_k \sigma\left(\frac{r_u - D_{ij}^k}{\tau}\right), \quad (2)$$

where σ is the sigmoid function, D_{ij}^k denotes the distance from pixel (i, j) to the k -th annotation point, and τ controls the softness of the confidence decline. Here the rate of the confidence bound decline is controlled by r_u and r_ℓ for the upper and lower bounds, respectively. Importantly, $r_u > r_\ell$. Example profiles of the two confidence bounds are shown in Fig. 2b. Regions above the upper bound or below the lower bound are disallowed regions, whereas regions between the two bounds are allowed ranges.

The *Window Loss* aims to discourage confidence predictions in any of the disallowed regions, while allowing any prediction within the allowed range. Dropping the pixel indexing for simplicity, we formulate the Window Loss as

$$WL(p, \ell, u) = \begin{cases} D(p, \ell) & \text{if } p \leq \ell, \\ 0 & \text{if } \ell < p \leq u, \\ D(p, u) & \text{if } p > u \end{cases}, \quad (3)$$

where $D(\cdot, \cdot)$ is a chosen divergence measure. Formulated this way, the Window Loss encouraged predictions violating the upper bound (resp., lower) to be lower (resp., higher), while not penalizing any predictions falling within the allowed regions. In term of divergence measures, we evaluate two options: mean squared error (MSE) and Kullback–Leibler divergence (KLD), written as

$$D_{MSE}(x, y) = \|x - y\|^2, \quad (4)$$

$$D_{KLD}(x, y) = y \log\left(\frac{y}{x}\right) + (1 - y) \log\left(\frac{1 - y}{1 - x}\right). \quad (5)$$

Examples of the Window Loss using MSE and KLD are shown in Fig. 2c. The Window Loss can garner significant improvements in performance where ambiguities are unavoidable. Its simple formulation allows it to be easily plugged into any existing deep learning framework.

Network Architecture

The radiographic appearances of pelvic fractures can vary considerably in both scale and shape. To represent the pathological patterns at different scales, we employ a ResNet-50

with a feature pyramid network (FPN) backbone (Lin et al. 2017a). The ResNet-50 portion has a bottom-up path that encodes the input PXR using four sequential blocks, with each down-sampling the image size by 2. The FPN adds a top-down path that upsamples and fuses the bottom-up feature maps. The feature maps are then processed by a 1x1 convolution to produce dense probability maps at four different scales for the input PXR image. The output probability maps can be interpreted as the confidence of presence of fractures at the corresponding locations in the PXR.

FPNs have been widely adopted in object detection to cope with the variance in object scales using multiple pyramid levels of feature maps with respect to different spatial resolutions. With bounding box annotations, objects are assigned to different pyramid levels according to their size. For fracture detection in PXR with point-based annotations, the size of the fracture is inherently ambiguous (from perspectives of both the complexity of the fractures and the form of annotation). Therefore, we assign every fracture annotation point to all FPN pyramid levels to encourage the network to recognize the fractures at different spatial resolutions. During training, losses are calculated on all pyramid levels. During inference, the probability maps from the four levels are ensembled and merged by averaging. The total training loss is the sum of the Window Loss on the four-level FPN outputs, written as:

$$L = \sum_{k=1,2,3,4} \mathbb{E}_{i,j} \left[WL\left(P_k(i, j), \ell_k(i, j), u_k(i, j)\right) \right], \quad (6)$$

where P_k is the probability map output at the k -th pyramid level, ℓ_k and u_k are the confidence bounds resized to the spatial resolution of P_k .

Implementation Details

We trained our model on a workstation with a single Intel Xeon E5-2650 v4 CPU @ 2.2 GHz, 128 GB RAM, 4 NVIDIA Quadro RTX 8000 GPUs. All methods are implemented in Python 3.7, and PyTorch v1.4. We use the ImageNet pre-trained weights to initialize the backbone network. The Adam optimizer with a learning rate of $4e - 5$, a weight decay of 0.001 and a batch size of 48 is used to train the model for 100 epochs. All images are padded to square and resized to 1024×1024 for network training and inference. We randomly perform rotation, horizontal flipping, intensity and contrast jittering to augment the training data. The supervision bounds are generated using parameters $r_\ell = 50$, $r_u = 200$ and $\tau = 2$. The trained model is evaluated on the validation set after every training epoch, and the one with the highest validation AUROC is selected as the best model for inference.

Experiments

Dataset: We retrieved 4410 PXR images of unique patients that were recorded from 2008 to 2016 in the trauma registry of Chang Gung Medical Hospital. Trauma-related findings are identified by a board consisting of a radiologist, a trauma surgeon, and an orthopedic surgeon who had fifteen years, seven years, and three years of experience, respectively. The

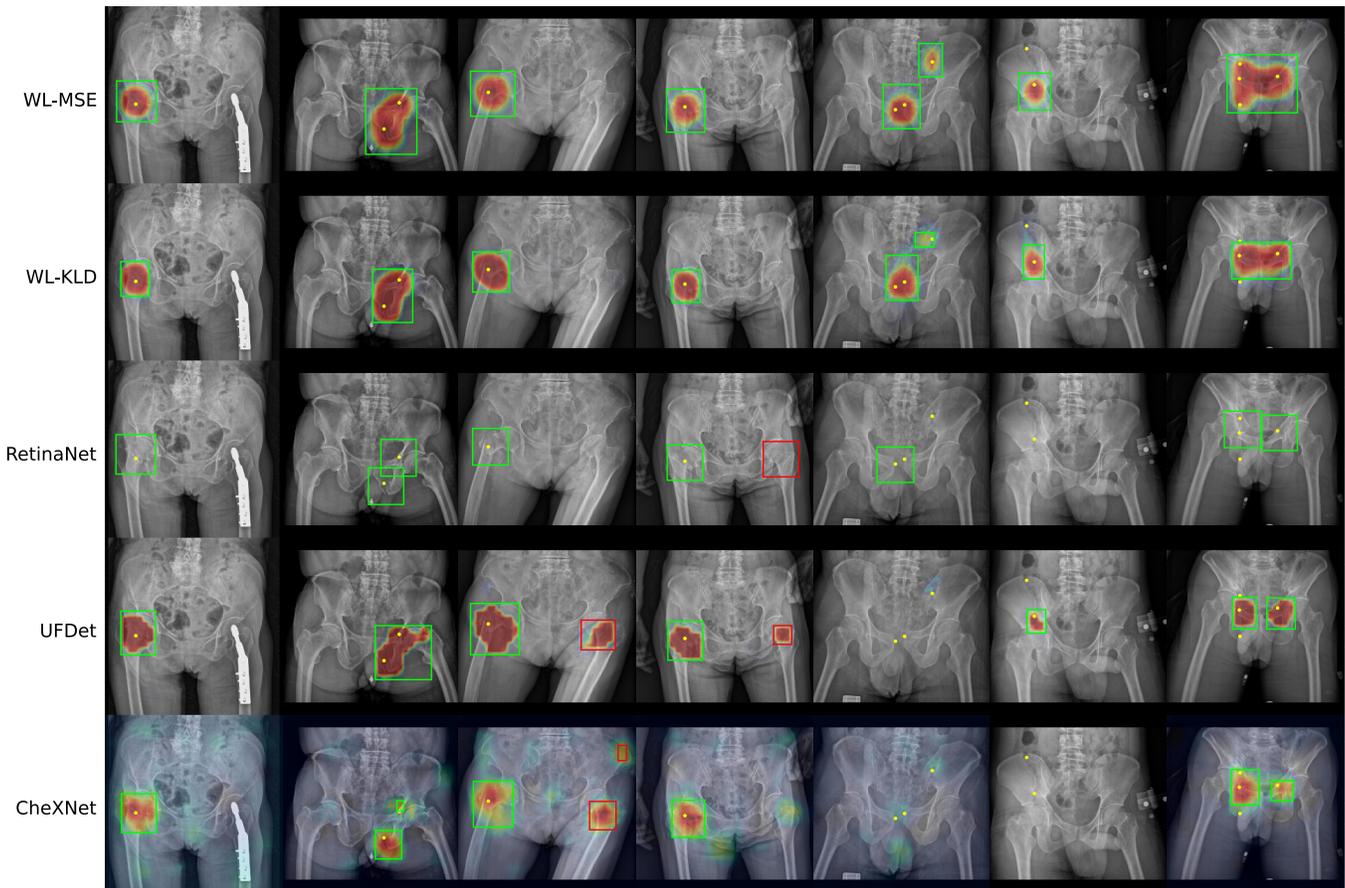


Figure 3: Visualization of fracture localization results. The proposed methods, WL-MSE and WL-KLD, achieve higher localization performances than the previous state-of-the-art methods. The localization performance of WL-MSE is qualitatively better than that of WL-KLD, which is in line with the FROC metrics.

best available information was provided, including the original images, radiologist reports, clinical diagnoses, surgical reports, and advanced imaging modality findings, if available. A total of 2776 images with acute trauma-related radiographic findings (1975 and 801 hip and pelvic fractures, respectively) are identified, resulting in 4066 annotated points (range 0-7). All experiments are conducted using five-fold cross-validation with a 70%/10%/20% training, validation, and testing split, respectively.

Evaluation Metrics

Classification metric We evaluate the image-level fracture classification performance using the receiver operating characteristic (ROC) curve and the widely used AU-ROC classification metric. For methods predicting probability map (e.g., the proposed method), the maximum response of the probability map is taken as the image-level classification score. For object detection methods predicting bounding box, the maximum classification score of all predicted boxes is taken as the image-level classification score.

Localization metric We evaluate the fracture localization performance of different methods using free-response oper-

ating characteristic (FROC) curve. To quantify FROC, we calculate an FROC score as the average recall at five false positive (FP) rates: (0.1, 0.2, 0.3, 0.4, 0.5) FPs per image. We also separately report the recall at FP=0.1 (Recall@0.1), the most clinically relevant setting. Given an operating point of the FROC, the predicted heatmap is binarized using the operating point to produce a mask. Connected component analysis is performed on the mask and each connected component is regarded as one detected fracture finding. A bounding box is then generated on each connected component. An annotation point is considered to be recalled if it is inside any bounding box. To define false positives, a ground truth mask is first generated with disks of radius 50 centered on the annotation points, which indicates areas that are certainly affected by the anomalies. A bounding box is considered as false positive if its intersection with the mask is less than 10% of its own region.

Comparison with Baseline Methods

Image classification methods We evaluate three state-of-the-art X-ray CAD methods based on image classification,

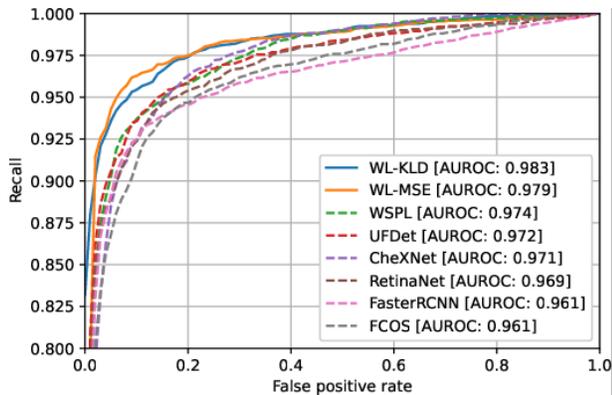


Figure 4: Comparison of receiver operating characteristic (ROC) curves (Recall scaled from 0.8 to 1.0).

CheXNet (Rajpurkar et al. 2017), *WSPL*¹ (Wang et al. 2017) and *UFDet*² (Wang et al. 2019). *CheXNet* trains a classification network with a GAP layer on the last feature map followed by a fully connected layer. *WSPL* replaces the GAP with a LSE pooling. *UFDet* estimates a dense probability map and use LSE pooling to produce the classification score. We evaluate the stage-1 of *UFDet* for a fair comparison with *CheXNet*, *WSPL* and our method, which are all single-stage methods. Anomaly localization map is generated using CAM (Zhou et al. 2016) for *CheXNet* and *WSPL*, and the produced probability map for *UFDet*. The localization map is converted to bounding box for FROC evaluation using the same steps described above. ResNet-50 is used as the backbone network for all image classification methods.

Object detection methods We evaluate three state-of-the-art object detection methods including two popular anchor-based detectors, *Faster-RCNN* (Ren et al. 2015) and *RetinaNet* (Lin et al. 2017b), and the latest anchor-less detector, *FCOS* (Tian et al. 2019). All compared methods use ResNet-50 with FPN as the backbone network. Since the actual scale of the fracture is unknown from the point-based annotation, a bounding box with a heuristically chosen size of 200×200 is placed on each annotation point. We empirically verified that the size of the bounding box is appropriate for the majority of fractures observed in PXR.

Result analysis Table 1 summarizes the results of image classification, object detection and the proposed methods. The proposed methods using both MSE and KLD as the divergence metric outperform both image classification and object detection methods on both the classification (*i.e.*, AUROC) and localization (*i.e.*, Recall@0.1 and FROC score) metrics by large margins. Specifically, WL-KLD achieves an AUROC of 0.983, outperforming the closest competitor, WSPL (0.974), by a health gap of 0.009. On Recall@0.1 and FROC score, WL-KLD leads the strongest baseline method, RetinaNet, by significant margins, *i.e.*, 2.8% (85.4% vs. 82.6%) and 1.5% (89.6% vs. 88.1%), respectively. WL-MSE

¹Weakly-Supervised Pathology Localization

²Universal Fracture Detection

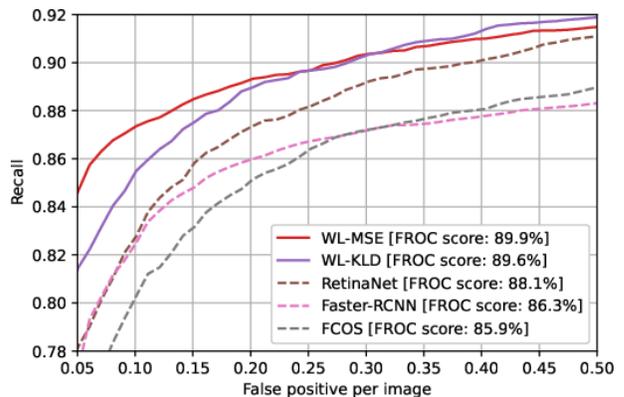


Figure 5: Comparison of free-response receiver operating characteristic (FROC) curves.

achieves even stronger localization performance than WL-KLD, reporting a Recall@0.1 of 87.4% and a FROC score of 89.9%. The AUROC of WL-MSE is lower than WL-KLD (97.9% vs. 98.3%), but it still outperforms all baseline methods by more than a 0.5%-1.8% margin.

An interesting observation is that although image classification methods do not use localization supervisory signals, they achieve superior classification performances comparing to object detection methods trained using localization supervision signals. In particular, the highest AUROCs achieved by image classification methods and object detection methods are 0.974 and 0.969, respectively, measuring a gap of 0.005. This suggests that algorithms designed specifically for detecting objects may not be optimal for recognizing the image-level existence of the object. On the other hand, image classification methods in general result in poor localization performance, with FROC scores between 54.9% and 75.4%, while all three object detection methods report substantially higher FROC scores between 85.9% and 88.1%. The high AUROC and low FROC score indicate that without localization supervision signals, the attention mechanism of image classification methods can activate wrong regions even when the image-level classification is correct. This is

Method	AUROC	Recall @FP=0.1	FROC Score
CheXNet	0.971	28.9%	54.9%
WSPL	0.974	56.0%	68.5%
UFDet	0.972	64.9%	75.4%
Faster-RCNN	0.961	82.5%	86.3%
RetinaNet	0.969	82.6%	88.1%
FCOS	0.961	80.3%	85.9%
WL-MSE	0.979	87.4%	89.9%
WL-KLD	0.983	85.4%	89.6%

Table 1: Fracture classification and localization results. Our method achieves the top performance, outperforming all baseline methods by significant margins.

τ	r_ℓ	r_u	AUROC	Recall@0.1	FROC Score
2	50	200	0.979	87.4%	89.9%
2	100	200	0.978	83.4%	86.4%
2	150	200	0.979	80.5%	82.0%
0.4	50	200	0.978	86.7%	89.3%
2	50	200	0.979	87.4%	89.9%
5	50	200	0.979	87.4%	88.2%
10	50	200	0.982	69.6%	84.3%
0	50	50	0.976	81.8%	86.1%
2	50	50	0.975	83.9%	87.2%

Table 2: Ablation study on the Window Loss with MSE.

also evidenced by the examples shown in Fig. 3.

While there is a trade-off between classification and localization performances using image classification and object detection methods, our method can simultaneously achieve improved performance in both aspects comparing to methods in both categories. Depending on whether MSE or KLD is used as the divergence measure in the Window Loss, our method behaves slightly differently in classification and localization, *i.e.*, KLD results in a slightly higher AUROC (98.3% vs. 97.9%), while MSE results in a slight higher FROC score (89.9% vs. 89.6%).

Ablation Experiments

We conduct ablation experiments to analyze the effects of the parameters of the Window Loss. We evaluate our method with varying parameters r_ℓ , r_u and τ using both MSE and KLD as the divergence measure. The results are recorded in Table 2 (using MSE) and Table 3 (using KLD).

We first evaluate two degenerate setting $[r_\ell, r_u = 50, \tau = 0]$ and $[r_\ell, r_u = 50, \tau = 2]$. In the first setting, the supervision bounds degenerate to a binary mask with disks of radius 50, and the Window Losses using MSE and KLD divergence measures degenerate to MSE and binary cross entropy (BCE), respectively. In the second setting, the supervision signal degenerate to a soft mask, similar to the Gaussian ground truth map widely used in key point localization methods (Wei et al. 2016). As shown in Table 2, both settings result in lower classification and localization performances comparing to our default setting $[r_\ell = 50, r_u = 200, \tau = 2]$. It demonstrates the effectiveness of the supervision bounds and the Window Loss. It is also worth noting that the degenerate setting, especially using KLD, still performs competitively comparing to image classification and object detection methods, measuring an AUROC score of 0.980 and a FROC score of 85.4%. It indicates the merit of modeling the anomaly detection task as a pixel-wise classification task.

To further analyze the effects of the parameters, we evaluate our method using $r_\ell = [50, 100, 150]$, $\tau = [0.4, 2, 5, 10]$ and $r_u = 200$. We use fixed $r_u = 200$ to simplify the analysis since it is empirically confirmed that areas more than 200 pixels away from any annotation points never have anomaly. As r_ℓ increases from 50 to 150, a bigger neighborhood of the

τ	r_ℓ	r_u	AUROC	Recall@0.1	FROC Score
2	50	200	0.983	85.4%	89.6%
2	100	200	0.984	81.2%	85.4%
2	150	200	0.984	82.8%	82.8%
0.4	50	200	0.982	85.5%	89.7%
2	50	200	0.983	85.4%	89.6%
5	50	200	0.984	82.0%	86.6%
10	50	200	0.986	69.2%	80.6%
0	50	50	0.980	78.3%	85.4%
2	50	50	0.981	83.0%	87.3%

Table 3: Ablation study on the Window Loss with KLD.

annotation point will have a higher lower bound, increasing the lower disallowed region. The results show that the classification performance remains very stable when r_ℓ increases, *i.e.*, the difference among the AUROC score is within 0.1. However, the localization performance degrades noticeably as the r increases, with the FROC score dropped by 7.9% (resp., 6.8%) using MSE (resp., KLD). We posit that because the larger positive region of a higher r_ℓ contains more normal areas, training the model to produce positive response in these regions can lower its localization accuracy. On the other hand, the results show that with a smaller $\tau = 0.4$, the classification performance remains almost the same as the $\tau = 2$ using both MSE and KLD, while the localization performance dropped slightly (0.7% in the FROC score) using MSE. When a larger $\tau = 5$ is used, the localization performance starts to degrade noticeably, *i.e.*, FROC score from 89.9% to 88.2% (resp., from 89.6% to 86.6%) using MSE (resp., KLD). As τ further increases to 10, the localization performance degrades significantly, *i.e.*, Recall@0.1 from 87.4% to 69.6% (resp., from 82.0% to 69.2%) using MSE (resp., KLD). This is because the overly smoothed supervision bounds have high tolerance of false activations in normal areas, preventing the training loss to provide sufficient guidance for localization.

Conclusion

In this paper, we demonstrate that anomalies or pathologies in X-ray imaging can have complex appearance and anatomical nature, which makes the bounding annotation unreliable. We present a disease detection method with a novel *Window Loss* that utilizes more flexible and efficient point-based annotations (capable of capturing anomaly findings in a nonparametric manner by point sampling). We validate the merit of our point-based annotation and Window Loss on the pelvic X-ray fracture detection task. The proposed method achieves substantially improved performance when comparing to state-of-the-art image classification and object detection methods. We also justify the usage of upper/lower bounds with allowed regions via our ablation study. Future studies should aim to employ the point-base annotation and Window Loss for detecting a broader range of anomalies in X-ray images from different body parts, as well as extending the method to new imaging modalities (*e.g.*, CT and MRI).

References

- Al-Masni, M. A.; Al-Antari, M. A.; Park, J.-M.; Gi, G.; Kim, T.-Y.; Rivera, P.; Valarezo, E.; Choi, M.-T.; Han, S.-M.; and Kim, T.-S. 2018. Simultaneous detection and classification of breast masses in digital mammograms via a deep learning YOLO-based CAD system. *Computer methods and programs in biomedicine* 157: 85–94.
- Baltruschat, I. M.; Nickisch, H.; Grass, M.; Knopp, T.; and Saalbach, A. 2019. Comparison of deep learning approaches for multi-label chest X-ray classification. *Scientific reports* 9(1): 1–10.
- Cai, J.; Harrison, A. P.; Zheng, Y.; Yan, K.; Huo, Y.; Xiao, J.; Yang, L.; and Lu, L. 2020. Lesion Harvester: Iteratively Mining Unlabeled Lesions and Hard-Negative Examples at Scale. *IEEE Transactions on Medical Imaging [Accepted]*.
- Chellam, W. 2016. Missed subtle fractures on the trauma-meeting digital projector. *Injury* 47(3): 674–676.
- Chen, H.; Wang, Y.; Zheng, K.; Li, W.; Chang, C.-T.; Harrison, A. P.; Xiao, J.; Hager, G. D.; Lu, L.; Liao, C.-H.; et al. 2020. Anatomy-aware siamese network: Exploiting semantic asymmetry for accurate pelvic fracture detection in x-ray images. In *European Conference on Computer Vision*, 239–255. Springer.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4700–4708.
- Huang, Y.-J.; Liu, W.; Wang, X.; Fang, Q.; Wang, R.; Wang, Y.; Chen, H.; Chen, H.; Meng, D.; and Wang, L. 2020. Rectifying Supporting Regions with Mixed and Active Supervision for Rib Fracture Recognition. *IEEE Transactions on Medical Imaging*.
- Irvin, J.; Rajpurkar, P.; Ko, M.; Yu, Y.; Ciurea-Illcus, S.; Chute, C.; Marklund, H.; Haghighi, B.; Ball, R.; Shpanskaya, K.; et al. 2019. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 590–597.
- Jiang, C.; Wang, S.; Liang, X.; Xu, H.; and Xiao, N. 2020. ElixirNet: Relation-Aware Network Architecture Adaptation for Medical Lesion Detection. In *AAAI*, 11093–11100.
- Law, H.; and Deng, J. 2018. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 734–750.
- Li, Z.; Wang, C.; Han, M.; Xue, Y.; Wei, W.; Li, L.-J.; and Fei-Fei, L. 2018. Thoracic disease identification and localization with limited supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8290–8299.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017a. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017b. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014a. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Lin, T.-Y.; Maire, M.; Belongie, S. J.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014b. Microsoft COCO: Common Objects in Context. In *ECCV*, 5:740–755.
- Liu, J.; Zhao, G.; Fei, Y.; Zhang, M.; Wang, Y.; and Yu, Y. 2019a. Align, attend and locate: Chest x-ray diagnosis via contrast induced attention network with limited supervision. In *Proceedings of the IEEE International Conference on Computer Vision*, 10632–10641.
- Liu, Y.; Zhou, Z.; Zhang, S.; Luo, L.; Zhang, Q.; Zhang, F.; Li, X.; Wang, Y.; and Yu, Y. 2019b. From Unilateral to Bilateral Learning: Detecting Mammogram Masses with Contrast Bilateral Network. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 477–485. Springer.
- Papadopoulos, D.; Uijlings, J.; Keller, F.; and Ferrari, V. 2017. Extreme clicking for efficient object annotation. In *ICCV*.
- Pham, H. H.; Le, T. T.; Ngo, D. T.; Tran, D. Q.; and Nguyen, H. Q. 2020. Interpreting Chest X-rays via CNNs that Exploit Hierarchical Disease Dependencies and Uncertainty Labels. *arXiv preprint arXiv:2005.12734*.
- Rajpurkar, P.; Irvin, J.; Zhu, K.; Yang, B.; Mehta, H.; Duan, T.; Ding, D.; Bagul, A.; Langlotz, C.; Shpanskaya, K.; et al. 2017. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, 91–99.
- Sirazitdinov, I.; Kholiavchenko, M.; Mustafae, T.; Yixuan, Y.; Kuleev, R.; and Ibragimov, B. 2019. Deep neural network ensemble for pneumonia localization from a large-scale chest x-ray database. *Computers & Electrical Engineering* 78: 388–399.
- Tian, Z.; Shen, C.; Chen, H.; and He, T. 2019. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE international conference on computer vision*, 9627–9636.
- Wang, X.; Peng, Y.; Lu, L.; Lu, Z.; Bagheri, M.; and Summers, R. M. 2017. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *IEEE CVPR*.
- Wang, Y.; Lu, L.; Cheng, C.-T.; Jin, D.; Harrison, A. P.; Xiao, J.; Liao, C.-H.; and Miao, S. 2019. Weakly Supervised Universal Fracture Detection in Pelvic X-Rays. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 459–467. Springer.

Wei, S.-E.; Ramakrishna, V.; Kanade, T.; and Sheikh, Y. 2016. Convolutional pose machines. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 4724–4732.

Yahalomi, E.; Chernofsky, M.; and Werman, M. 2019. 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*, 971–981. Springer.

Yan, C.; Yao, J.; Li, R.; Xu, Z.; and Huang, J. 2018a. Weakly supervised deep learning for thoracic disease classification and localization on chest x-rays. In *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, 103–110.

Yan, K.; Cai, J.; Harrison, A. P.; Jin, D.; Xiao, J.; and Lu, L. 2020. Universal Lesion Detection by Learning from Multiple Heterogeneously Labeled Datasets. *arXiv preprint arXiv:2005.13753* .

Yan, K.; Tang, Y.; Peng, Y.; Sandfort, V.; Bagheri, M.; Lu, Z.; and Summers, R. M. 2019. Mulan: Multitask universal lesion analysis network for joint lesion detection, tagging, and segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 194–202. Springer.

Yan, K.; Wang, X.; Lu, L.; and Summers, R. M. 2018b. DeepLesion: automated mining of large-scale lesion annotations and universal lesion detection with deep learning. *J. Med Imaging* 5(3).

Yao, L.; Prosky, J.; Poblenz, E.; Covington, B.; and Lyman, K. 2018. Weakly supervised medical diagnosis and localization from multiple resolutions. *arXiv preprint arXiv:1803.07703* .

Zhao, S.; Wu, X.; Chen, B.; and Li, S. 2019. Automatic Vertebrae Recognition from Arbitrary Spine MRI Images by a Hierarchical Self-calibration Detection Framework. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 316–325. Springer.

Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929.