

Breast mass detection with faster R-CNN: On the feasibility of learning from noisy annotations

Original

Breast mass detection with faster R-CNN: On the feasibility of learning from noisy annotations / Famouri, S., Morra, L., Mangia, L., Lamberti, F.. - In: IEEE ACCESS. - ISSN 2169-3536. - STAMPA. - 9:(2021), pp. 66163-66175. [10.1109/ACCESS.2021.3072997]

Availability:

This version is available at: 11583/2883374 since: 2021-11-24T17:37:51Z

Publisher:

IEEE

Published

DOI:10.1109/ACCESS.2021.3072997

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Received March 24, 2021, accepted March 31, 2021, date of publication April 13, 2021, date of current version May 7, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3072997

Breast Mass Detection With Faster R-CNN: On the Feasibility of Learning From Noisy Annotations

SINA FAMOURI¹, LIA MORRA¹, (Senior Member, IEEE), LEONARDO MANGIA,
AND FABRIZIO LAMBERTI¹, (Senior Member, IEEE)

Department of Computer and Control Engineering, Politecnico di Torino, 10129 Turin, Italy

Corresponding author: Sina Famouri (sina.famouri@polito.it)

This work was supported by NVIDIA Corporation through the donation of the Titan Xp GPU.

ABSTRACT In this work we study the impact of noise on the training of object detection networks for the medical domain, and how it can be mitigated by improving the training procedure. Annotating large medical datasets for training data-hungry deep learning models is expensive and time consuming. Leveraging information that is already collected in clinical practice, in the form of text reports, bookmarks or lesion measurements would substantially reduce this cost. Obtaining precise lesion bounding boxes through automatic mining procedures, however, is difficult. We provide here a quantitative evaluation of the effect of bounding box coordinate noise on the performance of Faster R-CNN object detection networks for breast mass detection. Varying degrees of noise are simulated by randomly modifying the bounding boxes: in our experiments, bounding boxes could be enlarged up to six times the original size. The noise is injected in the CBIS-DDSM collection, a well curated public mammography dataset for which accurate lesion location is available. We show how, due to an imperfect matching between the ground truth and the network bounding box proposals, the noise is propagated during training and reduces the ability of the network to correctly classify lesions from background. When using the standard Intersection over Union criterion, the area under the FROC curve decreases by up to 9%. A novel matching criterion is proposed to improve tolerance to noise.

INDEX TERMS Computer aided diagnosis, faster R-CNN, machine learning noise, mammography, object detection.

I. INTRODUCTION

In the last years, deep learning has led to major breakthroughs in many fields, including computer vision [1], medical imaging [2], [3], cyber-security [4], [5], and many others. In medical imaging, the performance of deep learning systems often surpasses that of conventional machine learning systems and in some cases even rival that of experienced physicians [3]. However, one of the crucial ingredients of this success is the availability of large curated image collections on which deep models can be trained. As a matter of fact, data starvation is often mentioned as one of the key obstacles to the application of deep learning in radiology [6]. Nonetheless, this notion is only partially correct [7]. Unlike other medical specialties, radiology departments are mostly digitized, and thousands if not millions of images

are routinely stored in picture archiving and communication systems (PACS). Besides privacy concerns, a key obstacle to exploitation is the high cost of collecting annotations [6], [7]. For the evaluation of Computer Aided Detection/Diagnosis (CAD) systems, a lesion-level reference standard (or ground truth) is usually established by marking lesions on the image based on an independent gold standard (e.g., biopsy) or, when not available, a panel of expert radiologists, to account for high inter-rater variability [8]. This strategy does not scale well to large datasets.

Information about lesion location and characteristics are routinely collected on PACS and reading workstations by radiologists, in the form of textual reports, various kinds of bookmarks, and lesion measurements [6], [9], [10]. This information could be harnessed to automatically mine an approximate reference standard with limited additional costs. However, compared to human-annotated datasets, quality is usually compromised for quantity, and it is important to

The associate editor coordinating the review of this manuscript and approving it for publication was Nilanjan Dey¹.

estimate the potential impact on the performance of deep neural networks trained on imperfect reference standard [9], [11], [12]. In the future, a more widespread adoption of structured reporting is expected to further facilitate this practice [6].

Some authors have used text mining to automatically extract labels from free-text reports [11]. Unfortunately, normally only patient-level labels can be generally obtained in this way; weakly supervised strategies such as Multiple Instance Learning (MIL) [13] can be used, but they usually perform poorly compared to fully supervised learning [14].

Recent works have shown the feasibility of establishing more precise and informative lesion-level annotations by mining hospital PACS [9]. Reporting workstations commonly provide drawing tools such as bounding boxes (ellipses or squares), arrows, lines or diameters, that radiologists can use to bookmark and measure specific lesions [9], [10]. A study conducted at the NIH Clinical Center found that the number of CT scans with such bookmarks skyrocketed after 2015; bookmarks often presented in the form of ellipses (8.4%) or lesion diameters (46%) [9]. The recently released DeepLesion dataset, which contains over 32,000 lesions identified on CT images based on diameter measurements, shows the potential of this approach [9]. Such annotations can be used to train object detection networks that, compared to image-level classifiers, can provide both lesion detection and localization [9], [15].

Mining strategies are attractive, but inevitably inject some levels of noise in the reference standard [14], [16]. For instance, there is no requirement that all the lesions mentioned in the report are explicitly annotated [9], and individual reports inevitably suffer from large inter-rater variability [17]–[19]. Research annotations are usually collected using two- or three-dimensional bounding boxes drawn as tight as possible to the lesion boundaries [10] or by segmenting the lesion [18]. Based on the authors' experience, bookmarks collected in clinical practice do not need to be as precise, and may serve additional purposes other than annotating the lesion (e.g., identifying the area selected for biopsy or further workup).

The present paper connects the problem of lesion detection from crowd-sourced annotations to the effect of noisy annotations on the generalization error of deep neural networks, and in particular object detection networks. In this context, previous studies have focused on changes in the ground truth labels (e.g., missing or mislabelled objects) [14]. Here, we identify a different and independent source of noise, which results from loosely annotated lesions (i.e., bounding boxes are approximate and include the lesion as well as part of the background). We show how this source of noise can potentially affect the detection performance, and investigate its implications by performing controlled experiments in which an increasing amount of noise is injected.

Many different architectures have been proposed for object detection [20]–[22] and successfully applied to medical imaging [15], [23]–[27]. However, all architectures share some common operating principles: they include one or

more classification modules that classify Regions of Interest (ROIs), each identified by a bounding box, as either background or one of the possible object (lesion) classes. These modules are trained by selecting examples of bounding boxes labeled as foreground (positive examples) or background (negative examples). Discrimination of positive and negative examples is achieved by comparing object proposals with the ground truth bounding boxes based on a *matching criterion*: usually, a threshold is set on the Intersection over Union (IoU) to determine if two boxes are a match. If the ground truth is imprecise (e.g., the bounding boxes are larger than the actual object), the matching may be incorrect, hence the classification modules will be trained on noisy labels and detection performance may suffer. It is precisely this phenomenon that we seek here to quantify and characterize.

Our experiments are conducted on the CBIS-DDSM (Curated Breast Imaging Subset of DDSM) dataset, a public screen-film mammography dataset for which high quality, clean labels are provided. In order to conduct a controlled experiment, noise is artificially injected by varying the size of the bounding boxes. Our reference architecture is the Faster R-CNN, which was shown to perform quite well for breast mass detection [15], [25]. It must be stressed, however, that matching object proposals with the ground truth is a common step for all object detectors and, while different detectors generate object proposals in different ways, matching is usually performed based on the IoU.

We establish the performance of the network in the presence of different noise levels and different matching criteria. When the bounding boxes increase in size, the number of ROIs labeled as positive increases, which is likely due to background being incorrectly labeled as foreground. Despite this fact, our experimental results show that object detection is quite robust in the presence of low to moderate amount of noise. In the presence of moderate to large noise, a simple yet effective countermeasure consists in the use of alternative matching criteria, favoring examples that are closer to the center of the ground truth bounding box and, thus, more likely to contain the actual lesion. Training lesion detectors by using imprecise bounding boxes is in principle feasible and, by carefully setting up the training procedure, even robust to high levels of noise.

The rest of the paper is organized as follows. Related works and background on Faster R-CNN are presented in Section II. Explanation of the noise model, network training, and the matching criteria used in this work are given in Section III. Experimental setup and results are reported in Section IV, and discussed in Section V.

II. RELATED WORK

A. OBJECT DETECTION AND FASTER R-CNN

Object detectors operate by classifying potential object regions that can be generated using fixed grids (one-stage detectors) or by employing a pre-selection mechanism (two-stage detectors). The Faster R-CNN [20] is a well-known two-stage architecture for object detection,

whereas single-stage architectures include YOLO [22] and RetinaNet [21].

Faster R-CNN is composed of two modules: the Region Proposal Network (RPN) and the detector. Both networks share the same convolutional backbone for feature extraction, which is usually pre-trained for classification and fine-tuned for object detection. Both the RPN and the detector have two heads, one for predicting the bounding box (regression head) and one for the classification.

The RPN takes an image (of any size) as input and outputs a set of object proposals or ROIs; it performs a binary classification to separate objects (of any class) from the background and narrow down the search. In order to do so, a sliding window is passed over the feature map and, at each location, k proposals with different scales and aspect ratios, known as *anchor boxes*, are generated (a common choice is $k = 9$). The RPN takes as input the coordinates of each anchor box, along with the feature map, and predicts the coordinates of the object bounding boxes, along with a binary score.

The RPN commonly predicts multiple overlapping bounding boxes for the same object. To further reduce the number of proposals, Non-Maximum Suppression (NMS) is used to reduce the number of proposals passed to the classifier: in short, for each group of overlapping bounding boxes, only the box with the highest classification score is retained, whereas the others are discarded. The detector then calculates the final classification score and the final bounding box, taking as input the feature map and the RPN proposals.

Training of the two heads is performed jointly in an alternating fashion. At each forward pass (where a forward pass corresponds to one image), the RPN is trained and updated; then, the output of the RPN is kept fixed and the detector head is updated. The loss is defined for both modules as a combination of a regression and classification loss [20]:

$$L(\{p_i\}, \{b_i\}) = \frac{1}{n_c} \sum_i L_{cls}(p_i, p'_i) + \lambda \frac{1}{n_r} \sum_i p_i * L_{reg}(b_i, b'_i) \quad (1)$$

where L_{reg} is the smooth L1 loss for regression, L_{cls} is the categorical cross entropy, b_i is the ground truth bounding box, b'_i are the output coordinates, p'_i is the predicted probability that b_i contains a lesion, and p_i is the ground truth label. This loss is used for both the RPN and the detector.

During training, each anchor box is labeled as a positive, negative or neutral example based on its overlap with the ground truth annotations. In the original Faster R-CNN architecture the overlap is calculated based on the IoU. Neutral examples, which are usually borderline (e.g., partially overlapping), are not used during training. The same process is repeated to train the detector. As it will be shown in Section III-C, the results critically depend on the definition of the matching criterion.

B. LABELING NOISE IN DEEP NEURAL NETWORKS

Real-world datasets are often affected by several forms of noise, which may prevent a machine learning model to correctly identify patterns in data. This is especially true in medical imaging, since accurate labels are expensive to obtain, pathological signs are often ambiguous, and inter-observer variability is high [16], [17]. The notion of label noise is not conclusive in the literature, where the term has been used to refer to different forms of label imperfections or corruption [16].

To the best of our knowledge, the effect of noise in training object detectors is relatively unexplored in literature, especially in the medical domain; compared to the general problem of object detection, lesions have much more ill-defined margins and are much rarer, leading to an extremely high-class imbalance. Therefore, the closest works in literature are those related to classification labelling noise, which has been extensively studied [14], [28]–[31]. Given a set of training samples $\{x_i, y_i\}$, where y_i is a discrete variable that corresponds to the true class of the sample, labelling noise can be formally defined as a stochastic process which pollutes the labels that are passed to the learning algorithm [30]. As a result, the observed label may no longer represent the true class of an instance.

Many noise models have been proposed for classification problems. The most common transformations studied are label flips and outliers [32]. Label flips are referred to samples that have been given a wrong class label, whereas outliers are samples which do not belong to any of the classes in the training set. Another important distinction is whether the noise affects all classes uniformly, or whether the noise is statistically dependent on the class or the features [30]. For instance, in the medical domain, negative or borderline results may be affected by higher noise since biopsy is not performed. We do not consider here the case in which label noise is constructed adversarially.

Various works have sought to establish the effect of labelling noise on classification performance in various regimes, either from a theoretical or experimental standpoint, and often reaching diverging conclusions. On the one hand, theoretical results [33] imply that a high capacity model should be robust to several types of random noise, provided that a sufficient number of training samples is available. To some extent, deep learning methods have indeed shown resilience to labeling noise given a required number of clean labels. For example, experiments on MNIST in [34] showed that with a 10:1 ratio of noisy to clean labels, at least 2000 clean labels are needed to reach an accuracy of 90%, whereas for a ratio of 50:1, the number of clean labels required increase to 10,000 for reaching the same performance. However, in practice the number of training samples is usually limited, especially for medical applications, and the label noise may follow complex, class-dependent patterns. Due to their memorization effect, sooner or later deep neural networks start to memorize noisy labeled samples, especially if the percentage of noisy labels is relatively large [35], [36].

It is therefore of practical interest to investigate experimentally the effect of noise in different tasks/datasets, in order to establish the feasibility of training deep neural networks [19], [37]. To test the robustness of algorithms against noise under controlled conditions, a common strategy is to inject synthetically generated noise into a clean gold standard dataset, progressively increasing the percentage of noisy labels [32], [37], [38]. While the most common strategy is to randomly flip a fraction of the labels, either based uniform or class-dependent random distribution, the generated noise is not always realistic, and other problem-specific strategies may be useful. For instance, Xue *et al.* simulated labelling noise in skin lesion classification by training a DNN on a fraction of the dataset, sorting the remaining images according to the test loss, and finally flipping the top $x\%$ of the high loss samples for both classes [37]. This method accounts for the fact that mislabelling by dermatologists does not depend on the class, but rather on the training sample, and is most likely to occur with difficult or borderline cases. They demonstrated a 20% decrease in accuracy when the noise ratio increased up to 40% [37].

In our work, we focus on object detection networks, and specifically on box coordinate noise (which we will refer to as *noise* in the remainder of this paper), assuming that the class label is *per se* correct. To the best of our knowledge, this type of noise has not been investigated yet. As we will show in our experiments, coordinate noise can translate into class label noise when training the RPN and detector classifier heads. In fact, as introduced in Section II-A, the training procedure labels anchor boxes as positive, negative, or neutral examples using a matching criterion. We will show how box noise can introduce labelling noise as a mix of class label flips (e.g., a positive anchor box becomes negative or viceversa) or outliers (e.g., a neutral bounding box becomes positive or negative).

Besides studying the practical effect of labelling noise, many authors have explored methods to make training *robust* in the presence of noise [28], [37], either by suggesting noise-tolerant losses [29], [32], [39], [40], or by detecting and removing noise from the dataset through dedicated networks or models [32], [36], [38]. The interested reader is referred to recent surveys for a more in-depth discussion on this subject [16].

In the field of object detection, previous work by Gao and colleagues showed how mining procedures, which are used to complete a partially annotated dataset, may introduce several types of labelling noise, such as false negatives, false positives and box coordinate noise [14]. To counter the effect of each type of noise, they added an ensemble classification head and a distillation head [41] in order to avoid overfitting to the labeling noise, and modified the loss to compensate the effect of false negatives. However, they trained the box regression heads only on seed annotations to circumvent the effect of box coordinate noise. Our work is different since we concentrate specifically on this type of noise, and show that it can affect the training not only of the regression, but also

of the classification heads. We also show that the matching criterion plays an important role in making the object detector more robust against noise, and propose ways to mitigate performance degrading.

C. BREAST MASS DETECTION

CNN-based models are currently the state of the art in breast image analysis, in some cases reaching performance closes to those of human experts [15], [23], [42]–[44].

Deep learning-based approaches to breast lesion detection and localization fall into two main approaches based on whether case-level or lesion-level annotations are available. At the case-level, recent works have shown that multi-stream CNNs can estimate the presence or absence of malignant findings from multi-view mammography images [23], [45]. Such experiments relied on large scale datasets, ranging from 100,000 to 200,000 cases and more than 1,000,000 images.

Other works have exploited CNN-based object detectors trained on lesion-level annotations, e.g., in the form of bounding boxes drawn around the lesion [15], [24]–[27], [44], [45]. From the radiologist perspective, this approach has the advantage of providing precise localization of cancerous lesions, which enables the method to be directly used as a CAD tool. By leveraging information about lesion localization during training, object detectors also have the potential to exploit training data more effectively [44]. Indeed, within the DREAM challenge [46], a Faster R-CNN model by Ribli *et al.* [15] reached the second position in the competitive phase, and scored first in the subsequent collaborative phase when more training data became available [44]. It should be noticed, however, that in the DREAM challenge itself, only case-level annotations were available, and the model had to be trained on separate private datasets, in part collected by the research group [44]. This fact highlights both the benefits and challenges of acquiring lesion-level annotations at large scale, since precise annotations have to be manually added by an experienced radiologist for research purposes. The present work aims at establishing whether these requirements can be relaxed, paving the way for future crowd-sourcing of annotations.

Several architectures were proposed for object detection. Compared to general-purpose object detection benchmarks, medical applications place a higher emphasis on accuracy than execution speed, as real-time performance is usually not needed. This is evident in mammography applications which employed architectures renowned for their accuracy, such as Faster R-CNN [15], [24], [25], [44] and RetinaNet [26], [45], with few exceptions based on YOLO [27]. On the public In Breast dataset [47], solutions based on Faster R-CNN outperformed those based on RetinaNet, achieving 92% sensitivity at 0.3 False Positive (FP)/image [24]. However, many differences in the experimental settings hinder direct comparison between papers: for instance, since the In Breast dataset is small, research groups relied on different public and private digital mammography datasets in order to build the training sets.

In our experiments, for the sake of reproducibility we exploited the publicly available CBIS-DDSM dataset, choosing the Faster R-CNN architecture based on the above considerations. Hyper-parameters of the Faster R-CNN architecture were tailored to the specific challenges of mammography, also taking into account previous works [15], [48]. We found in particular that Faster R-CNN overfits the CBIS-DDSM dataset, which is consistent with previous experiments by [25]. We also found that the choice of proposals during training has an important impact on generalization; thus, we propose to use hard sample mining on region proposals to increase performance (Section III-D).

D. METRICS AND LOSSES FOR BOUNDING BOX REGRESSION

A common problem of object detection and CAD systems is how to evaluate whether, and to what extent, an object proposal corresponds to a given ground truth bounding box. The IoU is the most common evaluation metric in object detection [20]–[22], [49]. Its value, bounded between 0 (no overlap) and 1 (complete overlap), measures how tightly the detector output fits the ground truth bounding boxes. As introduced in Section II-A, the quality of the bounding box depends mostly on the regression head.

Recently, two variants of the IoU have been proposed with the purpose of building more robust regression losses [49], [50]. Since the IoU is *per se* only defined when two bounding boxes overlap and does not provide any gradient for the non-overlapping cases, it cannot be used directly as loss for the regression head. Usually, a surrogate loss is used, such as the smooth l_1 -norm or, more generally, a l_p -norm [20]. The Generalized Intersection over Union (GIoU) [49] yields non-zero values for non-overlapping bounding boxes, thus overcoming one of the main limitations of the IoU. The Distance-IoU [50] combines the IoU with the Euclidean distance between the central points of the proposal and ground truth bounding boxes. Both metrics were used to define regression losses that could be incorporated in existing architectures, with performance gains especially for fast architectures such as YOLO [49], [50]. In the present research, the focus is on improving the performance of the classifier, rather than the regression head. Furthermore, our experimental settings assume that the bounding boxes in the ground truth have a systematic bias, which would be certainly captured by the regression parameters.

In this respect, the IoU metric still plays an important role, as it is used to label ROIs during training. In all detectors, ROI proposals that overlap a ground truth bounding box with an IoU greater than a certain threshold, such as 0.5 [21] or 0.7 [20], are marked as foreground, or true positives (TPs), whereas the rest are labeled as background, or false positives (FPs). Both [49] and [50] assume that the ground truth is noise-free, and do not investigate the properties of IoU, GIoU and Distance-IoU in the presence of noise.

When a bounding box proposal does not overlap with the ground truth, it can be safely labeled as background.

Hence, the GIoU metric [49] is not particularly relevant in this context, since the main difference with the IoU is in the case of non-overlapping bounding boxes. On the other hand, in the presence of noise the number of proposals that will overlap with the ground truth is expected to increase significantly, and matching criteria that incorporate both distance and overlap measures (e.g., [50]) may be able to better differentiate between competing overlapping proposals. In [50], however, this aspect was not investigated.

The present work also draws inspiration from the CAD evaluation field: here, compared to the field of object detection, a much wider range of mark-labeling or matching criterion have been proposed to decide which CAD marks correspond to the targeted abnormalities [8]. Mark-labeling rules are usually defined based on a measure of overlap, distance, or a combination them, also depending on the type of graphical marks used to represent the CAD output and/or to record the ground truth. Another common criterion is verifying whether the centroid of the CAD mark falls inside the ground truth bounding box [15]. Mark-labeling rules commonly used in object detection benchmarks may be too strict for CAD applications. For the clinical purposes of a CAD system, there is less interest on the tightness of the predicted bounding boxes, since lesions are sparse, often with ill-defined margins, and occlusions not existent. Previous research established that different criteria can result in dramatically different TP and FP estimates [8]. In this paper, different matching criteria are evaluated and compared in the presence of noise.

III. MATERIALS AND METHODS

The training framework for Faster R-CNN contains several modules that are summarized in Figure 1. Starting from the left, the input images alongside the annotations are passed to the network. The annotations are used to localize and identify the lesion and its type. Accordingly, during training, the annotations are considered as ground truth to calculate the loss and then via backpropagation the weights will be updated. However, if the bounding boxes are noisy (i.e., not tight to the lesions), we argue that they would lead to labeling noise when the region proposals are passed to the classifier, as discussed in Section III-C. Backpropagation with noisy labels will result in lower performance in lesion detection. Therefore, we set up a framework to analyse this issue and provide a solution to have more robustness against such noise. The rest of this section is dedicated to explaining each module in detail.

A. DATASET

The CBIS-DDSM collection [51]–[53] is an updated version of part of the DDSM dataset [54], selected and curated by a trained mammographer, and available for download from the Cancer Imaging Archive (TCIA). DDSM contains 2620 scanned screen film mammography studies, including normal, benign, and malignant findings with verified pathology information. Each study is composed of up to four

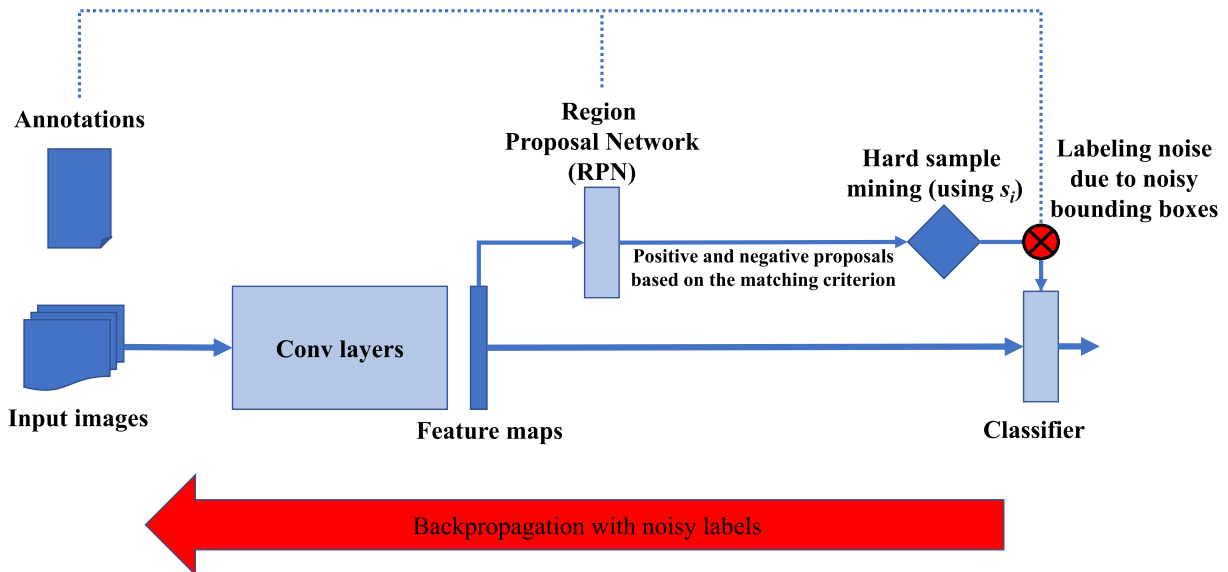


FIGURE 1. The training framework of Faster R-CNN is summarized in this figure, where the noise is injected on the bounding boxes in the annotations and its effect on the training is pinpointed by the red color. The red circle shows where the labeling noise is generated. Faster R-CNN is a two-stage object detector: the RPN filters candidate regions from the background, and the classifier assigns them the most likely class. Since lesions are rare compared to the background, a hard sampling procedure has been added to the framework to avoid overfitting and make sure that more informative proposals are passed to the classifier. If the reference bounding boxes are not tight to the lesions, the proposals that are passed to the classifier may contain noisy labels. In fact, proposals are automatically labelled against the reference standard on the basis of a matching criterion, such as the IoU, which may lead to incorrect results if the bounding boxes are not tight to the lesions.

images acquired in the cranio-caudal (CC) and medio-lateral oblique (MLO) orientations; however, only images with findings are included in the CBIS-DDSM dataset, for a total of 3,089 images. The CBIS-DDSM database has been pre-processed using standard techniques for screen-film mammography and converted to DICOM, as detailed in [52]. Additionally, we cropped the breast region before processing using a previously developed automatic algorithm [10], with the sole purpose of reducing computational time. In this study, we focus only on the detection of mass findings, for which accurate segmentation by experienced radiologists is available in the CBIS-DDSM collection. The ground truth bounding boxes are defined as the tightest bounding box that completely encloses the segmentation: we can therefore assume that the initial bounding boxes are very precise and essentially noise free. For microcalcification clusters, only a coarse segmentation is provided. Hence, we postulate that the type of noise we wish to study can be more accurately modelled for masses, and excluded such cases from the dataset. We used the standard training/test split (80%/20%) defined by the CBIS-DDSM authors. The final training and test set included 550 patients (with 613 masses) and 200 patients (with 222 masses), respectively. For each patient, up to four images are available and, since each finding is in most cases visible in both CC and MLO views, the total number of lesion views is 1,316 and 374 for the training and test set, respectively.

B. NOISE MODELING

Let us assume that a clean ground truth bounding box $b_i = (x_{1i}, y_{1i}, x_{2i}, y_{2i})$ is available for lesion $i \in \{1, 2, \dots, m\}$,

and that it perfectly fits the lesion boundaries. In order to inject noise, each b_i was modified by a random factor (noise), defined as:

$$w'_i = (1 + n_{wi})w_i, \quad h'_i = (1 + n_{hi})h_i. \quad (2)$$

where (w_i, h_i) are the width and height of b_i , (w'_i, h'_i) are the width and height of b'_i , and (n_{wi}, n_{hi}) are sampled from a normal distribution $n_{wi}, n_{hi} \sim \mathcal{N}(\mu, 1)$ with mean μ . Bounding boxes are defined in pixels in the ground truth, whereas μ is defined as a dimensionless multiplicative noise.

In practice, since the purpose of the bounding box would be to roughly bookmark the location of the region, we assume that the bounding box b'_i is always equal or larger than the clean bounding box, and that extremely large bounding boxes are likewise unlikely. Hence, we clipped n_{wi} and n_{hi} in the range $[0, 6]$, i.e., we assume the final bounding box is at most six times larger than the original.

Despite this limit, since the typical size of mammography masses ranges between 1 and 3 cm, the resulting bounding box may still exceed the breast region, which is unrealistic: based on the size of the largest lesions in CBIS-DDSM, we truncated b'_i to be at most 80% of the total image width. The center of b'_i is as same as b_i , except in those cases where the bounding box has been cropped to fit within these limits.

Four different levels of noise were generated with $\mu = \{0, 1, 2, 3\}$. Histograms of n_{wi} and n_{hi} drawn from the model are depicted in Figure 2, whereas the distributions of the bounding boxes diameters at different levels of noise are compared in Figure 3.




FIGURE 2. Histograms of the noise factor n_{wi} from $\mu = 0$ to $\mu = 3$. All distributions are clipped to the range $[0, 5]$.




FIGURE 3. Diameter distribution (histogram) of the noisy bounding boxes for the clean dataset and for each level of noise.

C. MATCHING CRITERION

As anticipated in Section II, a matching criterion is needed while training and testing to establish whether a bounding box proposal is a TP or FP detection. In other words, the matching criterion is a function that measures how well the ground truth bounding boxes, defined manually by experts, match the proposals given by the network. We included in the experiments three matching criteria, drawing from established papers in CAD evaluation and object detection: the IoU, which is the *de facto* standard in object detection, a simple Centroid-based criterion, and a combination of distance and overlap, which is denoted as Exp_IoU in the following.

Anchor boxes with a IoU higher than a threshold T_u are labeled as positive, and those lower than a second threshold T_l as negative. The rest are ignored at training time (neutral examples). In the rare cases in which no bounding box can be labeled as positive based on the IoU threshold, the anchor with the highest IoU overlap is selected [20]. As previously done in [15], we decrease the T_u from the original 0.7 value to 0.5 in order to allow more positive examples in each batch. We did not experiment with higher thresholds, as they were previously found to lead to unstable training [15]. T_l is instead equal to 0.3 [20].

The Centroid-based criterion or “centroid inside the bounding box” simply checks whether the center of the proposed bounding box falls inside the ground truth bounding box. This is a common criteria for evaluating CAD systems, but has never been used for training [15].

Overlap-based and distance-based metrics may produce different labels for a given bounding box. A simple example is given in Figure 4, which compares the ranking of proposed bounding boxes with respect to different criteria. Each criterion has different strengths: the IoU favors bounding boxes that match in both size and position, but may fail in the presence of noise (e.g., it may fail to match bounding boxes that are correct, but smaller, as in Figure 4A). Relying solely on the centroid is too weak, an intuition which is confirmed by experimental results in Section IV.




FIGURE 4. The red boxes show three different bounding box proposals for a lesion. If IoU is considered as the matching criterion, the order of the scores would be $C > B > A$. However, based on the distance between centroids, the order would be $A > C > B$.

Hence, we propose a new criterion, denoted Exp_IoU, which explicitly considers both the size and relative position of the bounding boxes, as follows:

$$S_{exp_iou}(b'_i, b_j) = \frac{IoU(b'_i, b_j) + e^{-\beta D(b'_i, b_j)}}{2} \quad (3)$$

where b'_i is the proposed bounding/anchor box, b_j corresponds to the ground truth, $D(\cdot)$ represents the Euclidean distance between the centers of each bounding box, and β balances the two contributions. The value of β was set to 0.1 based on experimental results. The same thresholds T_u and T_l are used for both IoU and Exp_IoU.

D. DEEP NETWORK ARCHITECTURE

As said, the architecture is Faster R-CNN, which was introduced in Section II-A. We here explain how the network was parameterized, and present some enhancements that were introduced to cope with the unique characteristics of the medical domain.

For the backbone network we chose the ResNet50 architecture, since preliminary investigations reported better results than VGG16 [15]. More specifically, layers up to conv4_x are included in the backbone, and the top layers (conv5_x) are included in the RPN and classifier heads.

The anchor box scales are $\{128, 256, 512\}$, whereas the aspect ratios are $\{(1, 1), (0.7, 1.4), (1.4, 0.7)\}$, with a stride of 16. Therefore, at each location nine anchor boxes will be generated. The NMS overlap threshold is set to 0.7 for training and 0.1 for testing; this threshold is much lower than in the original paper, but is justified by the fact that lesion margins in mammography are much less defined than object boundaries in natural scenes [15]. NMS limits the number of generated bounding boxes at 300 at both training and testing

time. The value of λ is set to 8.3 for the RPN loss, and 12.5 for the detector.

The RPN generates at most 300 examples per image. This results in a severely imbalanced classification problem, since the background is even more predominant than in general object detection. In fact, even considering only images which contain at least one lesion, the number of positive samples generated by the RPN is usually between 2 and 10. A possible strategy is to randomly sample a balanced subset: in our experiments, we chose four random ROIs (two positives and two negatives). This however may result in a poor selection of samples, and thus impair the network training.

A simple heuristic was added to mine informative examples after NMS. We assume that hard samples are those misclassified by the RPN, since one of the main difficulties in mammography is separating lesions from glandular patterns that mimic their presence. The following score is used to rank all the proposed ROIs:

$$s_i = (p'_i - p_i)^2 \quad (4)$$

where p'_i is the predicted probability that b_i contains a lesion, and p_i is the ground truth label. As s_i increases, the margin between the probability and the true label grows, meaning that it is a hard sample. In the presence of labeling noise, the margin s_i is also expected to be higher for noisy samples than clean ones [14].

The positive and negative samples are sorted separately based on their score s_i , and the corresponding mean scores S_P and S_N are calculated. The samples can then be split into four categories:

- Easy positive: positive samples with $s_i < S_P$;
- Hard positive: positive samples with $s_i \geq S_P$;
- Easy negative: negative samples with $s_i < S_N$;
- Hard negative: negative samples with $s_i \geq S_N$.

From each category, 25 positive samples and 25 negative samples are selected in order to maintain a balance between difficult and easy examples, as well as noisy and clean ones. Finally, four ROIs are randomly sampled from this subset.

E. EXPERIMENTAL SETUP

We performed a preliminary experiment in which we compared the baseline model (IoU matching criterion) with and without the hard negative sampling strategy. Then, performance across all levels of noise and matching criteria were compared, for a total of 15 possible configurations. All the experiments used the same hyper-parameters as explained in Section III-D. The thresholds for IoU and the Exp_IoU are defined in Section III-C. Weights pre-trained on ImageNet were used to initialize the backbone.

The Adam [55] optimizer was used with a learning rate of 10^{-5} ; each network was trained for 80 epochs, each comprising 500 iterations. Experimentally, we observed that Faster R-CNN starts to overfit after this number of epochs. Images were downsampled so that the largest dimension was equal to 600 pixels to reduce the computational effort,

even though if better results could be obtained with higher resolution images [15].

The network was implemented in Keras 2.2 with Tensorflow 1.13.1. In order to reduce the variability between different experiments, a constant seed was set for all libraries (Numpy and Tensorflow). The order of the images was randomized, but fixed for all experiments. All experiments were conducted on an Nvidia Titan Xp GPU.

F. EVALUATION

All experiments were evaluated on the clean test set, without inserting noise. The Free-Response ROC (FROC) was used to plot sensitivity vs. the average number of false positives (FPs)/image [8], [56]. The FROC paradigm is a location-specific variant of ROC analysis where the number of detections per image is not constrained, and each detection can be assigned a separate score by the CAD algorithm. The FROC curve is thus more suited to evaluate object detection networks, and is a widely accepted methodology in the medical image analysis literature [8]. The centroid inside the bounding box was used as the matching criterion for evaluation purposes. This is a commonly used methodology for evaluating CAD systems, and has been employed in similar works [15]. In fact, lesions are more sparse and have much less well-defined boundaries than objects on traditional images, and therefore tightly matching the reference bounding box is less relevant as long as the lesion is clearly shown to the radiologist. Furthermore, this choice allows us to compare performance at different noise levels on fair grounds. In fact, when comparing the predictions of networks trained on noisy bounding boxes against the clean ground truth, the IoU would significantly drop as the predicted bounding box is larger than the true lesion. Instead, we focus on whether each network can correctly locate the true lesion.

The FROC curves were computed on 1000 bootstrap samples, each containing 200 cases (which is the size of the validation set). All FROC curves were cut at 2 FPs/image, as high false positive rates are clinically less useful. The area under the FROC (AFROC) was used as a summary measure to compare experiments [56]. Confidence intervals were also calculated by bootstrapping.

IV. RESULTS

A. HARD SAMPLE MINING

Introducing hard sample mining improves the AFROC from 1.03 to 1.17 (+0.14), as depicted in Figure 5. At 0.5 FPs/image, this corresponds roughly to a 0.15 average increase in sensitivity. Therefore, this procedure was used in all further experiments.

B. EFFECT OF MATCHING CRITERIA AND NOISE ON BOUNDING BOX LABELING

In order to achieve a more in-depth understanding of the matching criteria and their tolerance towards noise, we analyzed the average number of anchors per lesion that were labelled as positive during the first iteration of the RPN training. Given that the actual number of lesions is constant,




FIGURE 5. Comparison of the FROC curves with and without applying hard sample mining.

we assume that an increase in the number of positive anchors corresponds to labelling noise due to the imperfect ground truth and/or imperfect matching criterion; in fact, as shown in Figure 6, many FP ROIs overlap with noisy ground truth boxes and may be incorrectly labelled. The results are shown in Figure 7 for both the clean and noisy datasets. A difference between the number of positive anchors is noticeable across different matching criteria, especially for the centroid criterion which is the most loose. It is also worth noticing that the amount of box coordinate noise results in a higher amount of labelling noise while training the RPN and the detector. The number of positive anchors increased up to eight times for the IoU and up to 10 times for the centroid-based criterion. On the contrary, for the proposed Exp_IoU criterion, the increase is only two-fold, and hence we hypothesize that the final performance will be less sensitive to noise.

C. FROC PERFORMANCE

The FROC curves for different noise levels and matching criteria are shown in Figure 8. For an easier comparison, we report the mean AFROC with respect to the noise level in Figure 9 and the corresponding confidence intervals in Table 1. The results confirm that the Centroid-based criterion is a poor choice for training, since the AFROC is always lower. The best results are achieved by both the IoU and the Exp_IoU on the clean dataset; however, for the IoU criterion the performance degrades almost linearly from 1.17 to 1.06 with increasing levels of noise. Exp_IoU is the most robust criterion with respect to noise, as no performance drop can be noticed.

By comparing Figure 7 and Figure 9, it can be noticed that the number of positive anchors and the performance are negatively correlated, which we attribute to a noisy ground truth degrading the performance of the network. Similar trends have been observed in classification networks where the amount of clean labels is relative small and the ratio of noisy to clean labels exceeds 10:1 [34].

V. DISCUSSION

In this work, we analyzed a specific type of noise (bounding box coordinate noise), which is extremely relevant to training

TABLE 1. Confidence intervals for the AFROC.

Criterion	Clean	$\mu = 0$	$\mu = 1$	$\mu = 2$	$\mu = 3$
IoU	1.29-1.05	1.29-1.09	1.23-1.039	1.23-1.01	1.16-0.95
Centroid	1.23-0.9	1.22-0.98	1.23-1.00	1.20-0.94	1.00-0.76
Exp_IoU	1.30-1.04	1.27-1.07	1.29-1.07	1.28-1.08	1.28-1.08




FIGURE 6. Clean (left) and noisy (right) ground truth box compared with ROI proposals generated by the RPN (shown in green). Many FPs overlap with the noisy ground truth box and, hence, may be incorrectly labeled as foreground.

object detectors from imperfect ground truth and was not covered by the previous literature. Our experiments show that coordinate noise, specifically in the form of *enlarged ground truth bounding boxes*, will result in labeling noise while training the classifier heads and, hence ultimately decrease the network performance.

The network tolerance to noise is mediated by the matching criterion used to label anchor boxes and RPN proposals during training. The standard criterion used in object detection, the IoU, is surprisingly robust to moderate levels of noise ($\mu = 1$) and degrades only for very high levels of noise: for instance, at 2 FPs/image the average sensitivity drops from 77% to 65%. Other criteria, such as in particular the Exp_IoU, are able to maintain performance stable across all levels of noise. The main reason is that it explicitly considers the position of the anchor box with respect to the lesion center, and not only the overlapping area. Therefore, misleading anchor boxes that are not well positioned will be discarded during training. However, this advantage may be lost if the lesion centers are affected by random noise as well.

Previous works had already shown that training object detectors for mammography, and medical images in general, requires specific choices of hyper-parameters, such as appropriate IoU thresholds, since lesions are less frequent and with more ill-defined borders compared to objects in natural scenes [15]. We confirmed such findings and employed a hard negative mining strategy to select informative and balanced samples for training the detector.




FIGURE 7. Average number of anchors per lesion labeled as positive in the first iteration of the RPN training. Results are compared for the Exp_IoU, IoU and Centroid-based criteria for the clean dataset (blue) and for increasing levels of noise (from yellow to purple). The number of positive anchors (and hence the noise) increases with more relaxed matching criteria and increases more than linearly with the amount of noise. All scales are logarithmic.




FIGURE 8. FROC curves with 95% confidence interval (calculated by bootstrapping). From left to right the IoU, Centroid, and Exp_IoU criterion were used. The latter is more tolerant towards noise with comparable performance across all levels of noise.

Our study has potential limitations. First of all, the dataset size is relatively small, and is not sufficient to reach state-of-the-art performance such as in [15]. This is partly due to overfitting, which was also observed in previous works [25]. In our experiments, overfitting was mitigated by early stopping, as well as by the hard sample mining procedure. Investigating to which extent overfitting is reduced with larger dataset sizes is an interesting avenue for future work. The impact of noise also depends on the size of the dataset. It is conceivable that, on larger datasets, the noise tolerance could further improve as more examples of lesions become available [34].

Secondly, with our current training procedure the effect of noise on the regression parameter is unavoidable, i.e., the predicted bounding boxes are going to be enlarged, as shown in Figure 10. We leave improving the quality of the regression to future work. For instance, knowing the level of noise in a given dataset, the regression parameters could be adjusted during or after training; strategies to mitigate the effect of bounding box noise on regression were also proposed by Gao and colleagues [14].

In this work, noise was simulated by explicitly manipulating annotations in a well-curated dataset. This is




FIGURE 9. Area under the FROC curves for the IoU (orange), the Centroid inside the bounding box (green) and the Exp_IoU criteria as a function of the noise level.

a common approach to study the effect of noise in machine learning and allows us to carefully control the experimental conditions. Nonetheless, our conclusions should be validated in a real-life dataset, which we plan on tackling in future work. Likewise, the proposed approach could be extended to other types of lesions or datasets.




FIGURE 10. Examples showing the clean and noisy ground truth annotations (red) vs. the network predictions (green) in the training set. From left to right the level of noise increases. The matching criterion used was the IoU. It can be seen that, as the noise level increases, additional false positives are produced by the network.

VI. CONCLUSION

We quantitatively investigated the effect of bounding box coordinate noise while training object detection networks

for mammography. We showed how state-of-the-art object detectors are robust to varying degrees of labelling noise, and proposed strategies to mitigate its effect at extreme

noise levels. Our study has important implications for dataset collection and annotation, since it shows that the bounding boxes do not need to be very precise for training to be effective. In the case of extreme noise levels, small changes in the training procedure, such as introducing a different matching criterion, can improve performance without increasing the complexity of the model, and can be easily incorporated in the training procedure of any object detector. These findings open new opportunities to train lesion detection model by using bookmarks and annotations routinely recorded by radiologists in their clinical practice.

REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- [2] L. Morra, S. Delsanto, and L. Correale, *Artificial Intelligence in Medical Imaging: From Theory to Clinical Practice*. Boca Raton, FL, USA: CRC Press, 2019.
- [3] B. Sahiner, A. Pezeshk, L. M. Hadjiiski, X. Wang, K. Drukker, K. H. Cha, R. M. Summers, and M. L. Giger, "Deep learning in medical imaging and radiation therapy," *Med. Phys.*, vol. 46, no. 1, pp. e1–e36, Jan. 2019.
- [4] Y. Xin, L. Kong, Z. Liu, Y. Chen, Y. Li, H. Zhu, M. Gao, H. Hou, and C. Wang, "Machine learning and deep learning methods for cybersecurity," *IEEE Access*, vol. 6, pp. 35365–35381, 2018.
- [5] M. Khosravy, K. Nakamura, N. Nitta, and N. Babaguchi, "Deep face recognizer privacy attack: Model inversion initialization by a deep generative adversarial data space discriminator," in *Proc. Asia-Pacific Signal Inf. Process. Assoc. Annu. Summit Conf. (APSIPA ASC)*, Dec. 2020, pp. 1400–1405.
- [6] M. D. Kohli, R. M. Summers, and J. R. Geis, "Medical image data and datasets in the era of machine learning: Whitepaper from the 2016 C-MIMI meeting dataset session," *J. Digit. Imag.*, vol. 30, no. 4, pp. 392–399, 2017.
- [7] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. W. M. van der Laak, B. van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, Dec. 2017.
- [8] N. Petrick, B. Sahiner, S. G. Armato, A. Bert, L. Correale, and S. Delsanto, "Evaluation of computer-aided detection and diagnosis systems," *Med. Phys.*, vol. 40, no. 8, 2013, Art. no. 087001.
- [9] K. Yan, X. Wang, L. Lu, and R. M. Summers, "Deeplesion: Automated mining of large-scale lesion annotations and universal lesion detection with deep learning," *J. Med. Imag.*, vol. 5, no. 3, 2018, Art. no. 036501.
- [10] L. Morra, D. Sacchetto, M. Durando, and S. Agliozzo, "Breast cancer: Computer-aided detection with digital breast tomosynthesis," *Radiology*, vol. 277, no. 1, pp. 56–63, 2015.
- [11] J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghighi, R. Ball, K. Shpankaya, J. Seekins, D. A. Mong, S. S. Halabi, J. K. Sandberg, R. Jones, D. B. Larson, C. P. Langlotz, B. N. Patel, M. P. Lungren, and A. Y. Ng, "CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison," 2019, *arXiv:1901.07031*. [Online]. Available: <http://arxiv.org/abs/1901.07031>
- [12] T. Tran and R. Kavuluru, "Distant supervision for treatment relation extraction by leveraging MeSH subheadings," *Artif. Intell. Med.*, vol. 98, pp. 18–26, Jul. 2019.
- [13] V. Cheplygina, M. de Bruijne, and J. P. W. Pluim, "Not-so-supervised: A survey of semi-supervised, multi-instance, and transfer learning in medical image analysis," *Med. Image Anal.*, vol. 54, pp. 280–296, May 2019.
- [14] J. Gao, J. Wang, S. Dai, L.-J. Li, and R. Nevatia, "NOTE-RCNN: NOise tolerant ensemble RCNN for semi-supervised object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9508–9517.
- [15] D. Ribli, A. Horváth, Z. Unger, P. Pollner, and I. Csabai, "Detecting and classifying lesions in mammograms with deep learning," *Sci. Rep.*, vol. 8, no. 1, p. 4165, Dec. 2018.
- [16] D. Karimi, H. Dou, S. K. Warfield, and A. Gholipour, "Deep learning with noisy labels: Exploring techniques and remedies in medical image analysis," *Med. Image Anal.*, vol. 65, Oct. 2020, Art. no. 101759.
- [17] D. Sacchetto, L. Morra, S. Agliozzo, and D. Bernardi, "Mammographic density: Comparison of visual assessment with fully automatic calculation on a multivendor dataset," *Eur. Radiol.*, vol. 26, no. 1, pp. 175–183, Jan. 2016.
- [18] S. G. Armato, G. McLennan, and L. Bidaut, "The lung image database consortium (LIDC) and image database resource initiative (IDRI): A completed reference database of lung nodules on CT scans," *Med. Phys.*, vol. 38, no. 2, pp. 915–931, Jan. 2011.
- [19] D. Kügler and A. Mukhopadhyay, "How bad is good enough: Noisy annotations for instrument pose estimation," 2018, *arXiv:1806.07836*. [Online]. Available: <http://arxiv.org/abs/1806.07836>
- [20] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 91–99.
- [21] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2980–2988.
- [22] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 779–788.
- [23] N. Wu, J. Phang, J. Park, Y. Shen, and Z. Huang, "Deep neural networks improve Radiologists' performance in breast cancer screening," *IEEE Trans. Med. Imag.*, vol. 39, no. 4, pp. 1184–1194, Apr. 2020.
- [24] R. Agarwal, O. Diaz, X. Lladó, M. H. Yap, and R. Martí, "Automatic mass detection in mammograms using deep convolutional neural networks," *J. Med. Imag.*, vol. 6, no. 3, 2019, Art. no. 031409.
- [25] K. H. Cha, N. Petrick, A. Pezeshk, C. G. Graff, D. Sharma, A. Badal, A. Badano, and B. Sahiner, "Reducing overfitting of a deep learning breast mass detection algorithm in mammography using synthetic images," *Proc. SPIE*, vol. 10950, 2019, Art. no. 1095004.
- [26] H. Jung, B. Kim, I. Lee, M. Yoo, J. Lee, S. Ham, O. Woo, and J. Kang, "Detection of masses in mammograms using a one-stage object detector based on a deep convolutional neural network," *PLoS ONE*, vol. 13, no. 9, Sep. 2018, Art. no. e0203355.
- [27] M. A. Al-masni, M. A. Al-antari, J.-M. Park, G. Gi, T.-Y. Kim, P. Rivera, E. Valarezo, M.-T. Choi, S.-M. Han, and T.-S. Kim, "Simultaneous detection and classification of breast masses in digital mammograms via a deep learning YOLO-based CAD system," *Comput. Methods Programs Biomed.*, vol. 157, pp. 85–94, Apr. 2018.
- [28] B. Frenay and M. Verleysen, "Classification in the presence of label noise: A survey," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 25, no. 5, pp. 845–869, May 2014.
- [29] S. Reed, H. Lee, D. Anguelov, C. Szegedy, D. Erhan, and A. Rabinovich, "Training deep neural networks on noisy labels with bootstrapping," 2014, *arXiv:1412.6596*. [Online]. Available: <http://arxiv.org/abs/1412.6596>
- [30] B. Frénay and A. Kabán, "A comprehensive introduction to label noise," in *Proc. ESANN*, 2014, pp. 1–10.
- [31] D. Hendrycks, M. Mazeika, D. Wilson, and K. Gimpel, "Using trusted data to train deep networks on labels corrupted by severe noise," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 10456–10465.
- [32] S. Sukhbaatar, J. Bruna, M. Paluri, L. Bourdev, and R. Fergus, "Training convolutional networks with noisy labels," 2014, *arXiv:1406.2080*. [Online]. Available: <http://arxiv.org/abs/1406.2080>
- [33] N. Natarajan, I. S. Dhillon, P. K. Ravikumar, and A. Tewari, "Learning with noisy labels," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 1196–1204.
- [34] D. Rolnick et al., "Deep learning is robust to massive label noise," 2017, *arXiv:1705.10694*. [Online]. Available: <https://arxiv.org/abs/1705.10694>
- [35] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning requires rethinking generalization," 2016, *arXiv:1611.03530*. [Online]. Available: <http://arxiv.org/abs/1611.03530>
- [36] B. Han, Q. Yao, X. Yu, G. Niu, M. Xu, W. Hu, I. Tsang, and A. Sugiyama, "Co-teaching: Robust training of deep neural networks with extremely noisy labels," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 8527–8537.
- [37] C. Xue, Q. Dou, X. Shi, H. Chen, and P.-A. Heng, "Robust learning at noisy labeled medical images: Applied to skin lesion classification," in *Proc. IEEE 16th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2019, pp. 1280–1283.
- [38] Y. Dgani, H. Greenspan, and J. Goldberger, "Training a neural network based on unreliable human annotation of medical images," in *Proc. IEEE 15th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2018, pp. 39–42.
- [39] A. Ghosh, H. Kumar, and P. Sastry, "Robust loss functions under label noise for deep neural networks," in *Proc. AAAI*, 2017, pp. 1919–1925.
- [40] G. Patrini, A. Rozza, A. K. Menon, R. Nock, and L. Qu, "Making deep neural networks robust to label noise: A loss correction approach," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1944–1952.

- [41] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015, *arXiv:1503.02531*. [Online]. Available: <http://arxiv.org/abs/1503.02531>
- [42] T. Kooi, G. Litjens, B. van Ginneken, A. Gubern-Mérida, C. I. Sánchez, R. Mann, A. den Heeten, and N. Karssemeijer, "Large scale deep learning for computer aided detection of mammographic lesions," *Med. Image Anal.*, vol. 35, pp. 303–312, Jan. 2017.
- [43] D. Lévy and A. Jain, "Breast mass classification from mammograms using deep convolutional neural networks," 2016, *arXiv:1612.00542*. [Online]. Available: <http://arxiv.org/abs/1612.00542>
- [44] T. Schaffter, D. S. M. Buist, C. I. Lee, Y. Nikulin, D. Ribli, Y. Guan, W. Lotter, Z. Jie, H. Du, and S. Wang, "Evaluation of combined artificial intelligence and radiologist assessment to interpret screening mammograms," *JAMA Netw. Open*, vol. 3, no. 3, 2020, Art. no. e200265.
- [45] S. M. McKinney, M. Sieniek, V. Godbole, J. Godwin, N. Antropova, H. Ashrafiyan, T. Back, M. Chesus, G. C. Corrado, and A. Darzi, "International evaluation of an AI system for breast cancer screening," *Nature*, vol. 577, no. 7788, pp. 89–94, 2020.
- [46] (2017). *Digital Mammography DREAM Challenge*. Accessed: Nov. 26, 2019. [Online]. Available: <https://www.synapse.org/#!/Synapse:syn4224222/wiki/434546>
- [47] I. C. Moreira, I. Amaral, I. Domingues, A. Cardoso, M. J. Cardoso, and J. S. Cardoso, "INbreast: Toward a full-field digital mammographic database," *Acad. Radiol.*, vol. 19, no. 2, pp. 236–248, 2012.
- [48] K. H. Cha, N. A. Petrick, A. X. Pezeshk, C. G. Graff, D. Sharma, A. Badal, and B. Sahiner, "Evaluation of data augmentation via synthetic images for improved breast mass detection on mammograms using deep learning," *J. Med. Imag.*, vol. 7, no. 1, 2019, Art. no. 012703.
- [49] H. Rezatofghi, N. Tsoi, J. Gwak, A. Sadeghian, I. Reid, and S. Savarese, "Generalized intersection over union: A metric and a loss for bounding box regression," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 658–666.
- [50] Z. Zheng, P. Wang, W. Liu, J. Li, R. Ye, and D. Ren, "Distance-IoU loss: Faster and better learning for bounding box regression," 2019, *arXiv:1911.08287*. [Online]. Available: <http://arxiv.org/abs/1911.08287>
- [51] R. S. Lee, F. Gimenez, A. Hoogi, K. K. Miyake, M. Gorovoy, and D. L. Rubin, "A curated mammography data set for use in computer-aided detection and diagnosis research," *Sci. Data*, vol. 4, no. 1, pp. 1–9, Dec. 2017.
- [52] R. S. Lee, F. J. Gimenez, A. Hoogi, and L. D. Rubin. (2016). *Curated breast imaging subset of DDSM [Dataset]*. *The Cancer Imaging Archive*. [Online]. Available: <https://doi.org/10.7937/K9/TCIA.2016.7002S9CY>
- [53] K. Clark, B. Vendt, K. Smith, J. Freymann, J. Kirby, P. Koppel, S. Moore, S. Phillips, D. Maffitt, M. Pringle, L. Tarbox, and F. Prior, "The cancer imaging archive (TCIA): Maintaining and operating a public information repository," *J. Digit. Imag.*, vol. 26, no. 6, pp. 1045–1057, Dec. 2013.
- [54] M. Heath, K. Bowyer, D. Kopans, R. Moore, and W. PhilipKegelmeyer, "The digital database for screening mammography," in *Proc. 5th Int. Workshop Digit. Mammography*, 2000, pp. 212–218.
- [55] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [56] A. I. Bandos, H. E. Rockette, T. Song, and D. Gur, "Area under the free-response ROC curve (FROC) and a related summary index," *Biometrics*, vol. 65, no. 1, pp. 247–256, Mar. 2009.



SINA FAMOURI received the master's degree in artificial intelligence from Shiraz University, Shiraz, Iran. He is currently pursuing the Ph.D. degree with the Dipartimento di Automatica e Informatica, Politecnico di Torino, Italy, supervised by Prof. Fabrizio Lamberti. His research interests include medical image analysis, computer vision, deep learning, and machine learning.



LIA MORRA (Senior Member, IEEE) received the Ph.D. degree in computer engineering from the Politecnico di Torino, Italy, in 2006. She is currently a Senior Postdoctoral Fellow with the Dipartimento di Automatica e Informatica, Politecnico di Torino. Her research interests include computer vision, pattern recognition, and machine learning. She is also an Associated Editor of *IEEE Consumer Electronics Magazine*.



LEONARDO MANGIA received the degree in biomedical engineering from the Politecnico di Torino, Italy, in 2019, with a focus on designing medical devices. He is currently working as a Biomedical Engineer in developing countries to improve their healthcare systems and their life-conditions. His research interests include machine learning and computer vision for medical applications.



FABRIZIO LAMBERTI (Senior Member, IEEE) is currently a Full Professor with the Dipartimento di Automatica e Informatica, Politecnico di Torino, Turin, Italy. His research interests include computer graphics, computer vision, human-machine interaction, and intelligent computing. He is also serving as an Associate Editor for *IEEE TRANSACTIONS ON COMPUTERS*, *IEEE TRANSACTIONS ON LEARNING TECHNOLOGIES*, *IEEE TRANSACTIONS ON CONSUMER ELECTRONICS*, *IEEE Consumer Electronics Magazine*, and the *International Journal of Human-Computer Studies*.