

SSMD: Semi-Supervised Medical Image Detection with Adaptive Consistency and Heterogeneous Perturbation

Hong-Yu Zhou^{a,b,*}, Chengdi Wang^{a,*}, Haofeng Li^{c,*}, Gang Wang^a, Shu Zhang^d,
Weimin Li^{a,†}, Yizhou Yu^{b,†}

^a Department of Respiratory and Critical Care Medicine, National Clinical Research Center for Geriatrics, Frontiers Science Center for Disease-related Molecular Network, West China Hospital, Sichuan University, Chengdu 610041, P.R. China

^b Department of Computer Science, The University of Hong Kong, Pokfulam, Hong Kong

^c Shenzhen Research Institute of Big Data, The Chinese University of Hong Kong, Shenzhen 518000, P.R. China

^d AI Lab, Deepwise Healthcare, Beijing 100080, P.R. China

Abstract

Semi-Supervised classification and segmentation methods have been widely investigated in medical image analysis. Both approaches can improve the performance of fully-supervised methods with additional unlabeled data. However, as a fundamental task, semi-supervised object detection has not gained enough attention in the field of medical image analysis. In this paper, we propose a novel Semi-Supervised Medical image Detector (SSMD). The motivation behind SSMD is to provide free yet effective supervision for unlabeled data, by regularizing the predictions at each position to be consistent. To achieve the above idea, we develop a novel adaptive consistency cost function to regularize different components in the predictions. Moreover, we introduce heterogeneous perturbation strategies that work in both feature space and image space, so that the proposed detector is promising to produce powerful image representations and robust predictions. Extensive experimental results show that the proposed SSMD achieves the state-of-the-art performance at a wide range of settings. We also demonstrate the strength of each proposed module with comprehensive ablation studies.

1 Introduction

Recently, deep convolution neural networks have achieved remarkable success in processing and understanding visual data, the convolution layers with learnable parameters in CNNs can adaptively harvest powerful image representations based on a large manually-annotated dataset. Some-

*First three author contributed equally.

†Corresponding author.

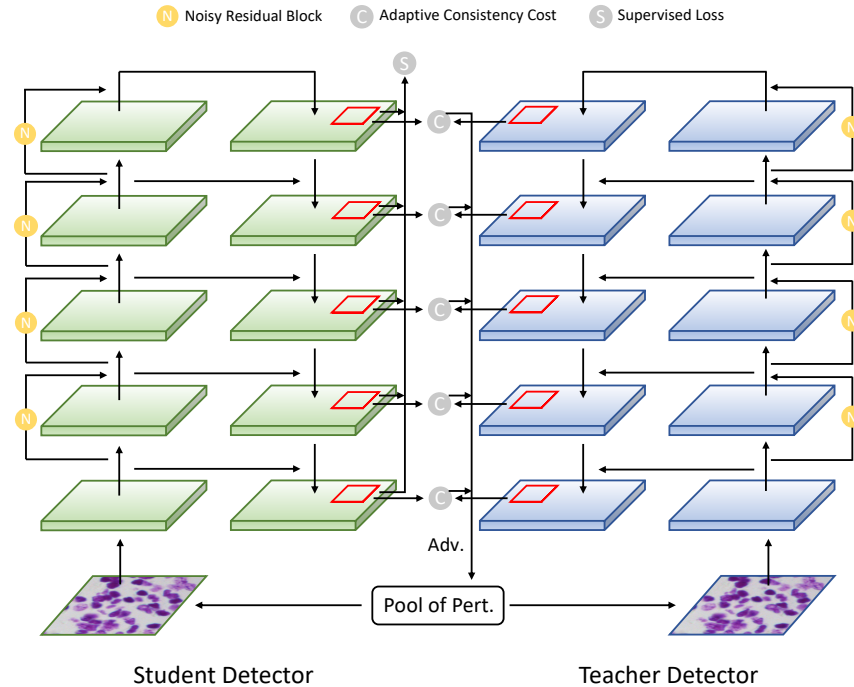


Figure 1: Overview of the proposed Semi-Supervised Medical image Detector (SSMD). Two feature pyramid networks are utilized to predict consistent outputs. **Pool of Pert.** refers to a set of perturbation strategies, which include horizontal flip, vertical flip, random rotation and adversarial perturbation (denoted as **Adv.**). Note that the adversarial perturbation is only applied to the input of teacher network.

times collecting and labeling such a large-scale dataset can be expensive, time-consuming and unaffordable in real applications. Particularly, annotating medical images usually require well-trained experts with prior biomedical knowledge. On the other hand, it is easier to obtain unlabeled data. Thus, how to train neural networks with both labeled data and unlabeled ones becomes an important problem, and also refers to the setting of semi-supervised learning.

Nowadays, a lot of deep semi-supervised methods [3, 17, 30, 35, 19] have been developed. They are effective in leveraging unlabeled data to mitigate the dependence of deep learning models on large-scale annotated datasets. In the field of medical image analysis, numerous semi-supervised learning methods have been proposed for a wide-range of applications, such as abnormality classification [18], 2D image segmentation [20, 43] and 3D volume segmentation [6, 31, 41, 20]. However, the community seldomly investigates the object detection task for medical images from a perspective of semi-supervised learning. In this paper, we mainly focus on two object detection tasks:

lesions detection and nuclei detection. Locating lesions and abnormalities in CT scans is a primary object detection task for radiologists. They need to find out the location of lesions, and describe the related attributes in radiological reports. An automatic lesion detector could not only reduce the workload of radiologists, but also benefit the areas that have a shortage of experienced radiologists. DeepLesion, a representative lesion detection dataset, contains 32,120 CT slices and 32,735 annotated lesions. Another fundamental object detection task in medical image analysis is nuclei or cell detection, which helps measure quantitative information to better understand disease progression. 2018 Data Science Bowl introduced a nuclei detection dataset that consists of about 27,000 cells. To reduce the high cost of medical image annotations, it is important to train a robust detector with not only labeled but also unlabeled medical images. Thus, in this paper we propose a novel Semi-Supervised Medical image Detector (SSMD) that can make use of unlabeled medical images to produce robust representations in an effective way.

Consistency-based semi-supervised learning methods [14] mainly utilize self-supervision [46, 5, 44] and usually consist of two procedures. First, synthesize a pair of input images via some data augmentation strategies. Second, force the paired model's outputs to be consistent, and formulate such constraint as an additional loss to train neural networks. The assumption behind is that deep learning models should produce similar representations or predictions with these augmented inputs. Therefore, such consistency constraint can be applied to the model outputs corresponding to the augmented inputs, and serves as an additional supervision when the ground-truth annotations are not available. However, some existing semi-supervised object detectors only adopt simple data augmentations such as image flipping, and do not consider the confidence of proposals in the consistency-based loss of semi-supervised detection.

Our proposed semi-supervised detector addresses two problems when applying the consistency regularization to medical image detection: a) too many background proposals may dominate the training procedure and b) mediocre augmentation strategies (such as horizontal flip and translation) cannot well regularize visual representations. For issue a), we propose *adaptive consistency cost* to adaptively scale the loss values, where the scaling factor decays to zero as confidence of the background class increases. Intuitively, the proposed adaptive mechanism can automatically down-weight the influences of background proposals. As for problem b), we introduce a set of heterogeneous perturbation methods to advance the regularization effect of the consistency loss. The core idea behind is that we want the detector to capture the invariant representations as we apply

various perturbations. We believe these representations are more robust and thus more generalized, even under various real-world noise.

2 Related Work

Many methods have been proposed to solve semi-supervised learning (SSL) problems. Here we mainly focus on deep learning based approaches which are most related to our work. Recent studies in semi-supervised learning could be categorized into two groups: pseudo labels and consistency regularization. In this section, we review these two types of methods from three different aspects: natural image classification, medical image analysis and semi-supervised object detection.

2.1 SSL in Natural Image Classification

[19] is the first to introduce pseudo-labeling for deep neural networks. Pseudo-labeling first trains a model with labeled data, predicts classification probability on unlabeled data, and then annotate the class with highest probability as pseudo labels. Lastly, all the data with real or pseudo labels are utilized to retrain the classification model. To obtain more accurate pseudo labels, [3] proposed MixMatch which employs a sharpen function to average the predictions of stochastic augmented inputs. In MixMatch, Mixup [42] was utilized as a strong augmentation method that could increase the diversity of both labeled and unlabeled data. [2] improved MixMatch by tackling the distribution misalignment and introducing more augmentation strategies. [38] provided a theoretical explanation for pseudo-labeling.

Besides from pseudo-labeling methods, [17] proposed π -model which encourages consistent network outputs between two realizations of the same input stimulus. Such consistency works as a supervision for unlabeled images, and is easily incorporated into training loss. [35] developed Mean Teacher that takes an exponential moving average of model weights to obtain more accurate predictions for consistency regularization. Different from Mean Teacher, [15] decoupled the connection between student and teacher networks, and used another student network to substitute for the teacher model. [36] proposed Interpolation Consistency Training which enforces the prediction at an interpolation of unlabeled points to match with the interpolation of the predictions at those points.

The most related work to ours is Virtual Adversarial Training [30] which improves the robust-

ness of the conditional label distribution around each input data point against local perturbation. However, our proposed detector is not simply transferring adversarial loss to semi-supervised classification. First, our proposed method makes use of position information to synthesize more effective adversarial samples. Second, the proposed adversarial perturbation considers the influence of different instances. Furthermore, we explore more perturbation strategies in both feature space and image space.

2.2 SSL in Medical Image Analysis

SSL methods are commonly used in medical image analysis to address the lack of manually annotated data. [18] proposed a patch-based semi-supervised learning approach and applied it to the classification of diabetic retinopathy from fundus images. [28], [29] and [40] used generative adversarial networks to conduct semi-supervised classification in chest X-ray, cardiology and dermoscopy, respectively. Recently, [45] developed a collaborative learning method to jointly improve the performance of disease grading and lesion segmentation, via an attention-based semi-supervised learning mechanism. [27] exploited unlabeled data by modeling the relation consistency among different samples, rather than only enforcing individual consistency.

Apart from image classification problems, SSL has also been applied to medical image segmentation task. [6] presented a new multi-task attention-based segmentation framework by enforcing consistency regularization on reconstructed foreground and background. [41] and [21] introduced a student-teacher framework which employs prediction uncertainty to highlight reliable consistent predictions. A soft-label based semi-supervised segmentation approach was presented in [4] to improve the ventricle segmentation of 2D cine MR images.

However, these above works do not aim at the object detection problem for medical images. To fill in such a gap, in this paper we develop a novel semi-supervised medical detector that emphasizes the importance of producing consistent and robust predictions.

2.3 Semi-Supervised Object Detection

[14] presented a Consistency-based Semi-Supervised learning for Object Detection (CSD) that works well for both single-stage and two-stage detectors. Compared with the consistency regularization in semi-supervised classification, the proposed consistency constraint in CSD is applied

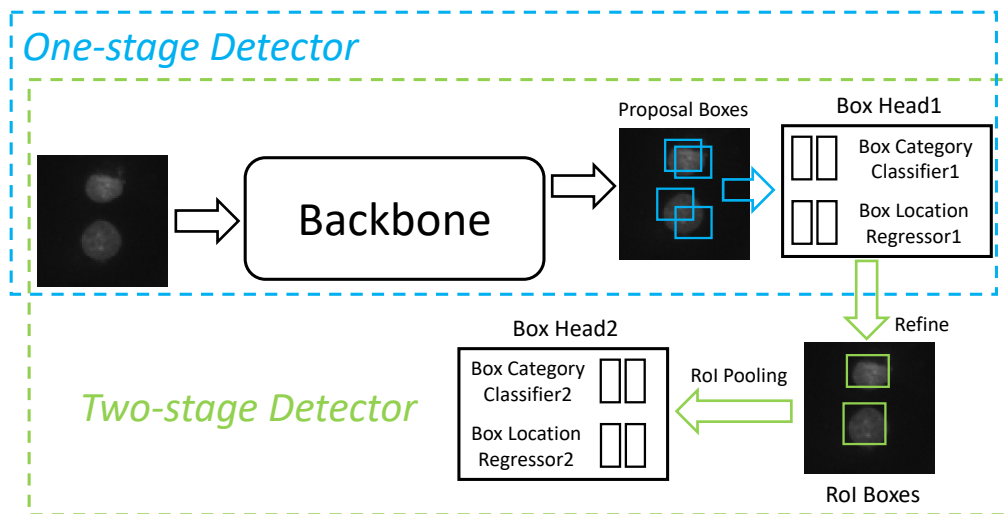


Figure 2: General network architecture for medical image detection. RoI is an abbreviation for region of interest.

to not only object classification but also object localization of a predicted region. Lately, following MixMatch [3], [37] built a novel semi-supervised lesion detector FocalMix based on Mixup [42] and a soft-version focal loss [25]. [33] proposed a semi-supervised learning object detection framework, STAC, which is based on high-confidence pseudo labels and the consistency via data augmentations. Different from STAC, our proposed method achieves different and stronger data augmentations by introducing Gaussian and adversarial noises to feature and image spaces.

In this paper, our proposed object detector is built on top of consistency regularization and CSD, instead of pseudo-labeling in FocalMix. Compared with CSD, the proposed approach employs heterogeneous perturbations to enhance the robustness of predictions as well as the detection accuracy. In addition, we develop a novel adaptive cost function to model instance-level consistency.

3 Method

Existing deep learning based medical image detection methods are usually anchor-based, which predict the relative position and scale factors between each object box and some pre-defined anchor boxes. These methods could be further divided into two types: one-stage [23, 16, 48] and two-stage detectors [37, 8, 47, 32, 26, 22], as shown in Figure 2. Given a one-stage detector, it first produces a large number of proposal boxes via the backbone network, after which one category classifier and

one location regressor are employed to deal with these boxes. In contrast, the two-stage pipeline requires one more box head which is responsible for refining those RoI boxes produced by the first box head. During the training stage, the overall loss function can be summarized as:

$$\text{CE}(p^c, p_{gt}^c) + \text{SmoothL1}(p_s^{\{x,y,w,h\}}, p_{gt}^{\{x,y,w,h\}}), \quad (1)$$

where p^c stands for the class prediction, p_{gt}^c denotes the ground truth class. Similarly, $p^{\{x,y,w,h\}}$ stands for coordinate and box size predictions while $p_{gt}^{\{x,y,w,h\}}$ represents their ground truth targets. CE stands for the cross entropy loss [10] to train box classifiers, and SmoothL1 represents the smooth L1 loss [10] which is employed to train box regressors. Recently, focal loss [25] is often used to replace the cross entropy loss when the number of proposals is extremely large [37, 48]. Both one-stage and two-stage approaches require backbone networks to extract image features, where 2D and 3D deep neural networks are used according to the types of input data. Specifically, for 2D tasks, VGG-16 [16] and ResNet [22] are two widely adopted architectures. For 3D tasks, 3D ResNet [37, 48] and 3D U-Net [23] are two representatives.

In this paper, we propose SSMD which incorporates medical image detection with semi-supervised learning. Compared to semi-supervised classification/segmentation, SSMD focuses more on instance regions instead of the whole image in classification or individual pixels in segmentation. Accordingly, to better regularize instance regions in detection, our SSMD addresses the importance of adding consistency to instance locations which are usually ignored in semi-supervised classification/segmentation.

In the following we describe three major contributions of SSMD: the adaptive consistency cost function, the noisy residual block and the instance-level adversarial perturbation strategy. We provide an overview in Fig. 1 in which a student-teacher framework is employed to generate predictions for shared inputs with different perturbation strategies. For labeled images, the proposed method uses an adaptive consistency cost and the supervised loss. For unlabeled data, only the adaptive consistency cost is used. The consistency loss is calculated with the predicted proposals at each spatial position and each scale.

To make use of unlabeled images, it is necessary to mine the data to generate intrinsic supervision signals which can be further incorporated into the training process. Nowadays, most semi-supervised deep learning approaches [17, 35, 30] focused on improving image classification results by keeping consistency within perturbed pairs. They require paired inputs where each pair contains the same image with different perturbation strategies. After feeding these pairs to neural

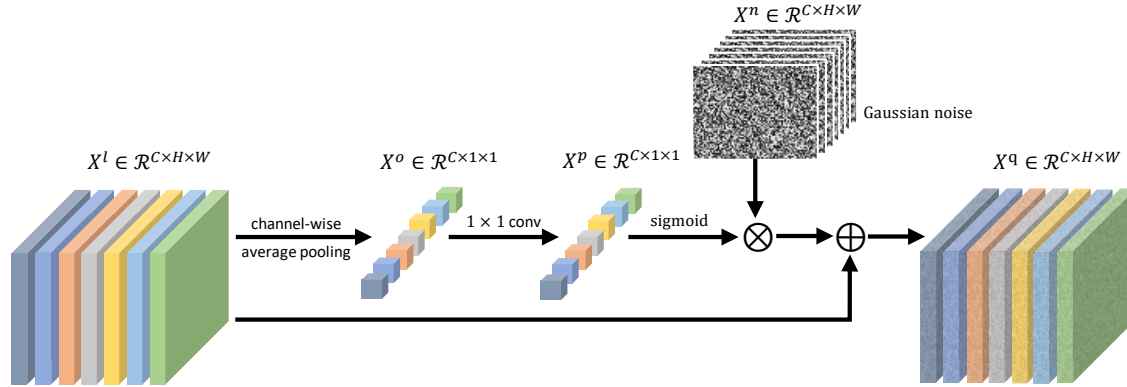


Figure 3: Illustration of the proposed noisy residual block. The proposed module adds noise perturbations to a residual block. Note that different colors mean different channels. \otimes stands for channel-wise multiplication while \oplus represents channel-wise addition.

networks, semi-supervised approaches force the outputs of each pair to be as close as possible. The most common perturbation methods can be summarized as: translation [17], rotation [35] and horizontal flip [17, 35, 30]. In this paper, we propose three more perturbation approaches: noisy residual block in feature space, instance-level adversarial perturbation and cutout in image space.

3.1 Adaptive Consistency Cost

As shown in Fig 1, the proposed SSMD model contains a student detector and a teacher detector where each network contains a feature pyramid network [24]. We adopt a parameter sharing approach proposed by [35] where the teacher model uses the exponential moving average (EMA) weights of the student model. We denote the weights of the teacher model and the student model as θ_t and θ_s respectively. θ_t^n denotes the weights of the teacher network at training step n and is updated as follows:

$$\theta_t^n = \alpha \theta_t^{n-1} + (1 - \alpha) \theta_s^n \quad (2)$$

where both θ_0^t and θ_0^s are independently initialized. During the training stage, for the student branch we apply random rotation and then randomly mask out some rectangle regions, which is known as cutout [7]. As for the teacher branch, we first apply horizontal flip and cutout to the augmented input of the student branch, and then add instance-level adversarial perturbation to it. Different from CSD [14], we propose to utilize an adaptive version of consistency cost to exploit unlabeled images and synthesize adversarial samples. The detector is based on RetinaNet which is to predict

the positions of proposals relative to pre-defined anchors. p^x, p^y, p^w, p^h , which are outputs of the proposed detector, denote four scale factors:

$$\begin{aligned} p^x &= (x - x^a)/w^a, p^y = (y - y^a)/h^a, \\ p^w &= \log(w/w^a), p^h = \log(h/h^a), \end{aligned} \quad (3)$$

where x, y are the coordinate of a proposal center. w and h represent the width and the height of a proposal. x^a, y^a, w^a and h^a are variables for default anchors. Let p^c denote the predicted probability distribution of different categories (after softmax). The whole procedure of the proposed semi-supervised medical detection is provided in Algorithm 1. We first apply different perturbations to a batch of labeled images \mathcal{X} for two branches, respectively. After the forward pass, we obtain the predictions of classes and box coordinates. For each labeled image in \mathcal{X} , its supervised loss (loss_{sup} , shown in Line 8 of Algorithm 1), which consists of a cross entropy loss (CE) and a smooth L1 loss (SmoothL1), can be directly calculated between the prediction and the ground truth.

To regularize the final predictions between the labeled images \mathcal{X} and the unlabeled images $\tilde{\mathcal{X}}$, we apply a consistency cost which includes a KL-Divergence loss (KL) and a mean squared error (MSE) loss, as shown in Line 11 of Algorithm 1. Assume that the output distributions of the teacher and the student models are close. Then KL loss is adopted for classification consistency, to measure the output difference between the teacher and the student networks. For location consistency, we follow the setting of CSD[14] and adopt MSE loss. Specifically, our proposed adaptive cost function contains a dynamic instance weight $W(p_s^c, p_t^c)$ which is defined as:

$$W(p_s^c, p_t^c) = \frac{(1 - p_s^c[0])^2 + (1 - p_t^c[0])^2}{2}, \quad (4)$$

where $p_s^c[0]$ refers to the probability belonging to the background category, predicted by the student network. $p_t^c[0]$ shares the same definition and is predicted by the teacher network. In our implementation, we treat the features of different levels equally in the adaptive consistency cost. For each feature level, the weight of the adaptive cost is equal to 1. The MSE loss displayed in Line 11 and Line 12 of Algorithm 1 is computed as:

$$\begin{aligned} \text{MSE}(p_s^{\{x,y,w,h\}}, p_t^{\{x,y,w,h\}}) &= \text{MSE}(p_s^x, p_t^x) + \text{MSE}(p_s^y, p_t^y) \\ &+ \text{MSE}(p_s^w, p_t^w) + \text{MSE}(p_s^h, p_t^h), \end{aligned} \quad (5)$$

where p_s and p_t are the predictions of student model and teacher model, respectively. For the prediction of unlabeled data \tilde{p} , we calculate its MSE loss $\text{MSE}(\tilde{p}_s^{\{x,y,w,h\}}, \tilde{p}_t^{\{x,y,w,h\}})$ in a similar

way with Equation (5). Note that during the inference stage, only the student network $f_{\theta_s^N}(\cdot)$ is used to perform final predictions.

The proposed adaptive consistency cost takes into account the predicted confidence of proposals at each spatial position. Given a proposal with high foreground probability, it would result in a higher weight of the consistency cost than those of easily recognized background samples. This mechanism helps the model apply more regularization effects to objects instead of the meaningless background. In practice, this adaptive cost is applicable to both labeled and unlabeled medical images, making proposed detector more effective in the setting of small amounts of labeled data.

3.2 Noisy Residual Block

In this part we propose *noisy residual block* that adds noise to intermediate feature maps. The proposed noisy residual block can be regarded as a perturbation strategy working in a feature space. As shown in Fig.3, we modify the classical residual block used in [13] and append an attention-based mechanism. We name the proposed module noisy residual block, since it introduces noise perturbations to a residual block. More details are in the following.

The input to layer l is denoted as $X^l \in \mathcal{R}^{C \times H \times W}$. The proposed noisy residual block first applies a channel-wise average pooling to X^l and then adopts a 1×1 convolutional operation:

$$X^p = conv(AvgPool(X^l)). \quad (6)$$

where $X^p \in \mathcal{R}^{C \times 1 \times 1}$ and *AvgPool* is the abbreviation of global average pooling. For each layer l , we sample a Gaussian noise map $X^n \in \mathcal{R}^{C \times H \times W}$ where each component is drawn from a Gaussian distribution $\mathcal{N}(\mu, \sigma)$. μ and σ stand for the mean and standard deviation, respectively. Meanwhile, we employ a scaled sigmoid function to normalize X^p . A channel-wise multiplication is performed between X^p and X^n . Finally, X^q can be computed by adding the multiplication result to the input feature X^l :

$$X^q = (X^n \otimes sigmoid(\gamma X^p)) \oplus X^l. \quad (7)$$

where γ is a scale factor. Here we employ a *sigmoid* function to adaptively control the noise level of different channels in the noise perturbation. X^q serves as the output of the noisy residual block and will be passed to following layers.

An intuitive understanding of the noisy residual block is to add “appropriate” noise to intermediate representations. For example, shallow layers are supposed to have wild noise as they are foundations of the whole network. We believe the degree of the embedded noise should be determined and can be learned by the representations themselves. Motivated by this idea, the noisy residual block learns channel-wise attentions to apply channel-dependent noise to feature maps. Moreover, we employ a residual connection to maintain the stability of the training process.

3.3 Instance-level Adversarial Perturbation based on Consistency Regularization

Adversarial training has been widely adopted as a useful way to improve semi-supervised classification and segmentation. In contrast, the detection problem focuses more on instances instead of pixels in classification or segmentation. Thus, the methods designed for classification/segmentation may not be suitable for detection because they treat all pixels equally. In this section, we propose an instance-level adversarial perturbation strategy to address this issue.

Let r_{adv} denote the adversarial perturbations added to the input image. In each training iteration, r_{adv} is first initialized from a normalized Gaussian distribution and has the same shape as \mathcal{X} and $\tilde{\mathcal{X}}$. Then, a scaled r_{adv} is added to the original image as:

$$\text{Adv.}(\mathcal{X}) = \mathcal{X} + \xi r_{adv}, \quad (8)$$

where ξ is a scale factor satisfying $0 < \xi \leq 1$. Classical adversarial examples work by causing classifiers to predict a wrong category. However, in SSMD, the goal of adding adversarial perturbations is to increase the difficulty of performing consistency regularization. Note that similar computation process can also be applied to $\tilde{\mathcal{X}}$.

We follow [11] to synthesize a well perturbed input. We pass $\{\mathcal{X}, \tilde{\mathcal{X}}\}$ and $\{\text{Adv.}(\mathcal{X}), \text{Adv.}(\tilde{\mathcal{X}})\}$ to student and teacher networks respectively, to obtain the consistency loss loss_{cont} (shown in Line 11 of Algorithm 1). Only the high-confidence predictions are used to compute the consistency loss for gradient backward when applying adversarial perturbation. The gradient g and the adversarial perturbations r_{adv} are computed as:

$$\begin{aligned} g &= \nabla_{r_{adv}} \text{loss}_{cont} * \mathbf{1}[\sum p_s^c \parallel \sum \tilde{p}_s^c > \tau], \\ r_{adv} &= \epsilon \frac{g}{\|g\|}. \end{aligned} \quad (9)$$

where the symbol \sum denotes the sum of all foreground classes. $\mathbb{1}[\cdot]$ is an indicator function which equals 1 when $\sum p_s^c$ or $\sum \tilde{p}_s^c$ is larger than a given threshold τ . ϵ is the strength of perturbation, controlling the magnitude of r_{adv} . $\|\cdot\|$ stands for L2 normalization. After computing Equation (9), r_{adv} is added to \mathcal{X} to obtain the final perturbed input. In general, it requires an additional forward and backward pass to synthesize the perturbed input image before we feed these final inputs to the detection network. Such process is to maximize the effect of r_{adv} on loss_{cont} , and can be viewed as an adversarial process.

Similar to the adaptive cost, we design instance-level perturbation to amplify the influences of high-confidence foreground proposals while reducing the impacts of low-confidence ones. In practice, foreground pixels receive heavy adversarial noise while the perturbation of background pixels has much smaller magnitude. Such implementation makes the consistency loss focus more on foreground objects, producing effectively perturbed inputs.

4 Experiments

In this section, we first conduct ablation studies to better understand the strengths of different modules in the proposed method SSMD. Moreover, we design comprehensive experiments to verify the effectiveness of SSMD on various settings.

4.1 Dataset

The experiments are conducted on a nuclei dataset and a lesion database. For both datasets, we manually and randomly split the training set into labeled data and unlabeled data with fixed ratios in order to fit the setting of semi-supervised learning.

Nuclei Dataset In our experiments, we adopt the nuclei dataset introduced by 2018 Data Science Bowl¹ (DSB, hosted by Kaggle). The dataset was acquired under a variety of conditions and includes nuclei images of different cell types, magnifications, and imaging modalities. The training set contains 522 nuclei images (80%) while the validation set has about 60 images (10%). The rest images are used for testing. On average, each image contains about 45 cells which are enough to train a robust nuclei detector. In practice, we only assign labels to some training images and take the other training images as unlabeled data. The evaluation metric is mAP

¹<https://www.kaggle.com/c/data-science-bowl-2018>

DeepLesion Dataset We also present experimental results on DeepLesion [39] which is a large-scale public dataset containing 32,120 axial Computed Tomography (CT) slices of 10,594 studies collected from 4,427 patients. The dataset has 32,735 annotated lesion instances in total. Each slice contains 1~3 lesions. The additional slices above and beneath a target slice are regarded as relevant contexts of the target slice. These additional slices are of 30 mm. In most cases, a slice is 1 or 5 mm thick. The dataset covers a wide scope of lesions from lung, liver, mediastinum (essentially lymph hubs), kidney, pelvis, bone, midsection and delicate tissue. Following [48, 22], we test our proposed method on official testing set (15%) and report the sensitivity at 4 false positives (FPs). We directly use the training and validation set officially provided by DeepLesion.

4.2 Implementation Details

For DSB dataset, the proposed detector is built on top of an ImageNet-pretrained ResNet-50 which has five scales. Nine default anchors are adopted in each scale. The size of input images is 448×448 . The batch size is 8. All models are trained for 100 epochs. Adam is utilized as the default optimizer with $1e-5$ as the initial learning rate, which is then divided by 10 at the 75th ($100 \times \frac{3}{4}$) epochs. For the supervised baseline, image rotation and horizontal image flipping are considered as default augmentation strategies. It is worth noting that the hyperparameter λ of consistency loss (shown in Line 13 of Algorithm 1) plays an important role during the training stage. We first gradually increase the value of λ to 1 in the first quarter of the training, and then decrease it to 0 in the last quarter. The formal definition of λ is:

$$\lambda = \begin{cases} e^{-5(1-\frac{4j}{N})^2}, & 0 \leq j < \frac{N}{4} \\ 1, & \frac{N}{4} \leq j < \frac{3N}{4} \\ e^{-12.5(1-\frac{7(N-j)}{N})^2}, & \frac{3N}{4} \leq j \leq N \end{cases} \quad (10)$$

where N is the number of training iterations and j is the iteration index. Similarly, for DeepLesion dataset we simply follow the preprocessing method in [48] to resize each slice into 512×512 pixels whose mean voxel-spacing is 0.802mm. We first clip the Hounsfield units (HU) to $[-1100, 1100]$ and then normalize them to $[-1, 1]$. We compute the mean and standard deviation of the whole training set and use them to further normalize input slices. For both datasets, we set γ to 0.9 and the degree of random rotation is set to 10 degrees.

Dataset	Method	10%	20%	30%	40%	50%
DSB	Supervised	39.3	47.3	53.1	58.9	62.3
	Pseudo-labeling	46.8	53.5	57.2	62.5	65.5
	CSD	45.7	54.0	58.5	62.9	65.2
	FocalMix	48.9	56.2	60.1	64.0	67.0
	SSMD	52.2 +3.3	59.1 +2.9	62.3 +2.2	65.9 +1.9	68.8 +1.8
DSB	Fully-supervised	76.7				
-	<i>p</i> -value	0.0154	0.0232	0.0195	0.0274	0.0242
DeepLesion	Supervised	48.0	54.9	61.6	67.3	71.5
	Pseudo-labeling	55.6	59.8	64.7	70.2	73.6
	CSD	55.3	59.5	65.2	71.0	74.0
	FocalMix	57.6	62.2	66.9	72.4	75.3
	SSMD	60.4 +2.8	64.1 +1.9	68.4 +1.5	74.1 +1.7	76.5 +1.2
DeepLesion	Fully-supervised	86.5				
-	<i>p</i> -value	0.0096	0.0104	0.0083	0.0074	0.0098

(a) Backbone: ResNet-50

Dataset	Method	10%	20%	30%	40%	50%
DSB	Supervised	42.1	50.6	55.7	61.0	64.5
	Pseudo-labeling	49.0	55.6	59.4	64.7	66.8
	CSD	48.2	55.8	59.9	65.1	66.2
	FocalMix	51.0	58.4	62.3	65.9	67.4
	SSMD	54.8 +3.8	61.7 +3.3	64.9 +2.6	68.2 +2.3	69.4 +2.0
DSB	Fully-supervised	78.3				
-	<i>p</i> -value	0.0121	0.0135	0.0093	0.0101	0.0089
DeepLesion	Supervised	50.8	57.3	64.0	69.5	73.0
	Pseudo-labeling	57.7	61.5	66.6	72.1	74.7
	CSD	57.1	61.8	67.4	73.2	75.1
	FocalMix	59.8	64.3	69.0	74.5	76.5
	SSMD	62.5 +2.7	66.4 +2.1	70.7 +1.7	76.1 +1.6	77.9 +1.4
DeepLesion	Fully-supervised	87.9				
-	<i>p</i> -value	0.0073	0.0082	0.0065	0.0061	0.0078

(b) Backbone: ResNet-101

Table 1: Comparison with baseline approaches under different labeled ratios. For all SSMD experiments, we apply cutout to their inputs. We report experimental results on both ResNet-50 and ResNet-101. All fully-supervised baselines make use of 100% labeled data. The input resolution of DSB is set to 448×448 while the input size of DeepLesion is 512×512 . The best results are in bold and we also display the relative improvements compared to the second best results. All *p*-values are computed between SSMD and FocalMix.

4.3 Baselines

Supervised and fully-supervised detectors For supervised detector, we only use the labeled data to train deep models. As for the fully-supervised baseline, we train our nuclei and lesion detector using the whole training set. In practice, for both two approaches, we save checkpoints based on

their performance on validation sets. These models are then used to perform the test whose results are reported in the following.

Pseudo-labeling An intuitive thought in semi-supervised classification for utilizing unlabeled data is to use a trained model to make predictions, which can also be applied to object detection. However, considering the fact that existing deep learning based detectors usually produce unreliable results given a small amount of training data, we have to cherry-pick optimal predictions from dozens of region candidates. It is a laborious and tedious process. Therefore the strategy used in [33] is applied to filter low-confidence candidates with a high confidence threshold. Pseudo labels with their corresponding images and labeled data are utilized to train a new detector from scratch.

Consistency-based Semi-Supervised learning method for object Detection (CSD) CSD [14] employs consistency constraints as a tool to improve detection performance by making full use of unlabeled data. For each input image, CSD first applies horizontal flip to construct an input pair which is fed to a siamese network to obtain two sets of predictions. Constraints are then added to regularize these predictions and serve as an additional supervision for unlabeled images. Compared to CSD, our proposed method emphasizes the robustness of predictions under various perturbations.

FocalMix FocalMix [37] was the the first approach to investigate the problem of semi-supervised learning for medical image detection. FocalMix is based on MixMatch [3] whose idea is similar to pseudo-labeling. Mixup [42] is used as the main augmentation strategy in labeled training set and a soft-target focal loss is proposed to leverage soft targets. The original version of FocalMix is a 3D detector while we extends it to a 2D version.

Implementation of baselines For fairness, all baselines share the same detection backbone (i.e., RetinaNet), input sizes and training strategies with those of our SSMD, if not specified otherwise. For Pseudo-labeling, we set the confidence threshold to 0.9 following the instruction from [33] and use one-shot encoding to represent pseudo-labels. For CSD, we use KL divergence and MSE loss to regularize detector's predictions. Since the code of FocalMix has not been released, we directly implement a 2D detector (based on RetinaNet) following the official paper[37]. That is to say, we employ the same mixup strategies including image-level mixup and object-level mixup to augment training data. For mixup hyperparameters, we set η to 0.2. We also implement the soft-target focal loss together with the sharpening operator, where α_0 is set to 0.05, α_1 is set to 0.95, γ is set to 2.0 and the sharpen temperature factor T is set to 0.7. We have tried to train FocalMix for {100,200,400,800} epochs, where we find that training for 200 epochs achieves the

best performance in both DSB and DeepLesion.

4.4 Comparison with the Baselines

The experimental results of supervised and semi-supervised approaches are reported in Table 1. We conduct experiments on both DSB and DeepLesion datasets with different network architectures and different labeled ratios.

The first observation is that the proposed SSMD can consistently outperform other semi-supervised baselines by a significant margin in different datasets and labeled ratios. Particularly, SSMD has greater advantage in small labeled ratios, such as 0.1 & 0.2. For example, SSMD outperforms FocalMix by more than 3 points in DSB when the labeled ratio is 10% while such gap becomes smaller as the labeled ratio increases. Similar phenomena can also be observed in DeepLesion dataset. We argue that the significance in small labeled ratios may be due to the fact that more perturbations lead to stronger regularization, preventing deep neural networks from the overfitting scenario which is pretty common given a small number of training data. Such argument is supported by the following observation: the performance gap between SSMD and other semi-supervised methods is larger in DSB than the gap in DeepLesion. Because DeepLesion is a large-scale dataset and even 10% of it has thousands of training samples which definitely help deep learning models avoid overfitting.

Another observation is that FocalMix consistently surpasses pseudo-labeling and CSD while pseudo-labeling performs as well as CSD in most cases. Typically, pseudo-labeling seems to obtain better results in small labeled ratios whereas CSD performs better in large ratios. Such trend is obvious in DeepLesion but becomes less visible in DSB. Nonetheless, all semi-supervised methods significantly outperform the supervised baseline, which verifies that semi-supervised methods are considerably helpful with small ratios of labeled data.

In Fig.4, we conduct experiments on different approaches with various input resolutions and two different network architectures. We can see that SSMD outperforms FocalMix significantly in all resolutions on both DSB and DeepLesion datasets. If we compare CSD with pseudo-labeling, we can see that although pseudo-labeling can outperform CSD on DSB dataset under small labeled ratios, it cannot beat CSD on DeepLesion dataset. For SSL with different network architectures, we can find that the gap between SSMD and other baselines becomes smaller when network becomes deeper (ResNet-50 vs. ResNet-101). We argue that deeper models may reduce the relative improvements of most of SSL algorithms on both DSB and DeepLesion datasets.

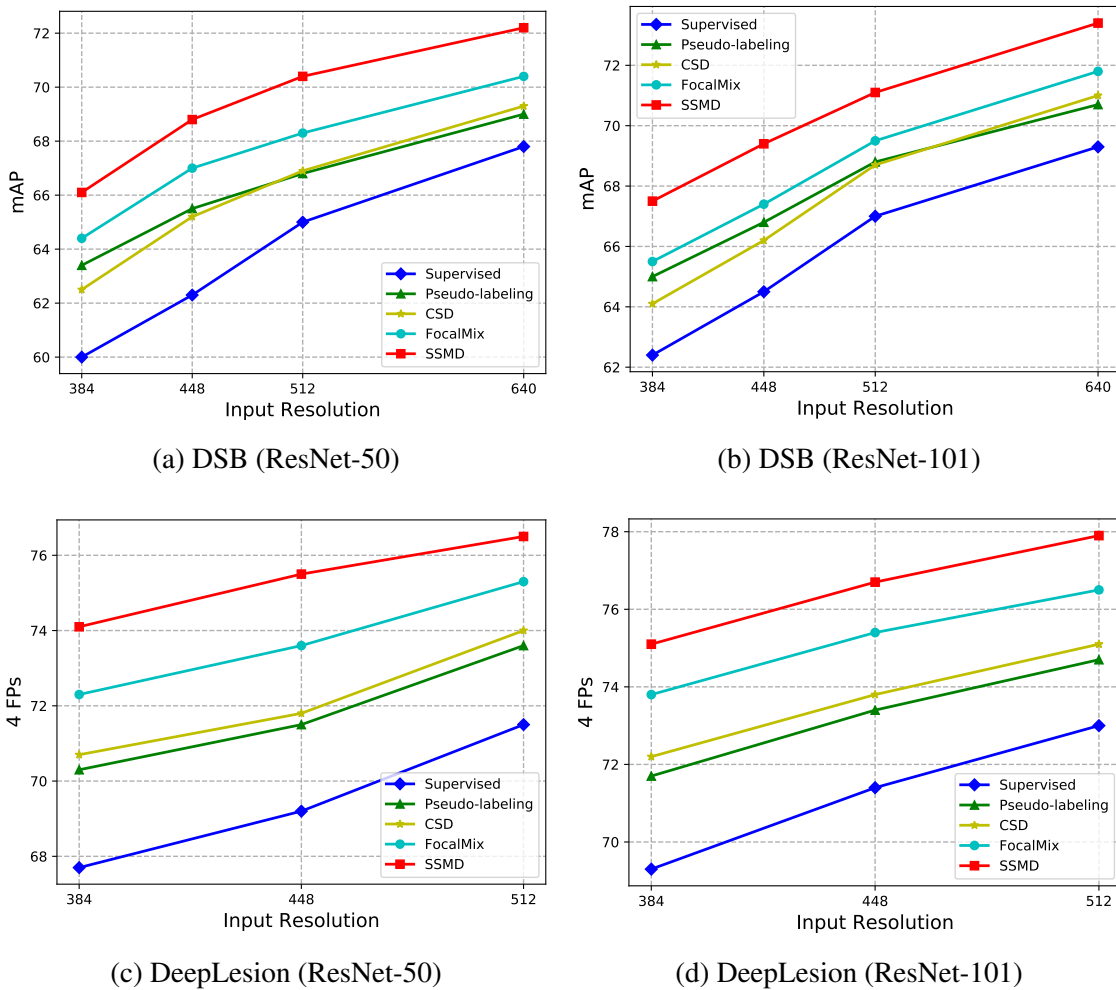


Figure 4: Comparison of different semi-supervised detectors with various input resolutions. mAP and 4 FPS are used for DSB and DeepLesion as evaluation metrics. The labeled ratio is 50%.

ACC	NRB	IAP	DSB	DeepLesion
			65.2	74.0
✓			66.7	75.1
✓	✓		68.0	75.9
✓	✓	✓	68.8	76.5
<i>p</i> -value			0.0227	0.0328

Table 2: An incremental study of all proposed modules. The default input resolutions for DSB and DeepLesion are 448×448 and 512×512 , respectively. **ACC** stands for adaptive consistency cost, **NRB** denotes the proposed noisy residual block and **IAP** represents the instance-level adversarial perturbation. The *p*-values are calculated between the top-2 models in each column.

Dataset	ACC	384	448	512	640
DSB		62.5	65.2	66.9	69.3
DSB	✓	63.8	66.7	68.6	71.1
<i>p</i> -value		0.0174	0.0133	0.0121	0.0107
DeepLesion		70.7	71.8	74.0	-
DeepLesion	✓	71.6	72.8	75.1	-
<i>p</i> -value		0.0154	0.0143	0.0127	-

Table 3: Influence of the proposed adaptive consistency cost (ACC). Note that for SSMD without the adaptive consistency cost, we directly employ the consistency loss used in CSD.

4.5 Ablation Study

In this section, we first make an incremental study on all proposed modules. Then, we study how different settings and factors affect the proposed SSMD. These factors include the adaptive consistency cost (ACC), the perturbation strategies in feature space, the hyperparameters of adding adversarial perturbations (ξ , ϵ and τ), the EMA factor α , different consistency constraints (c , x , y , w and h), different number and size of masks for applying cutout, and adding consistency loss to different numbers of feature scales. Considering that these hyperparameters could have different affects with different resolutions, we present the results with different input sizes. For all experiments, *the default labeled ratio is 0.5 and the network architecture is ResNet-50*. We repeat each

Dropout	Dropblock	Noise	NRB	DSB (mAP)				DeepLesion (4 FPs)		
				384	448	512	640	386	448	512
				62.5	65.2	66.9	69.3	70.7	71.8	74.0
✓				62.8	65.0	66.3	69.2	70.1	71.3	73.6
	✓			62.9	65.6	67.3	69.4	71.0	<u>72.3</u>	74.2
		✓		62.5	65.3	66.9	69.1	70.8	71.7	73.9
			✓	63.6	66.2	67.8	70.1	71.9	72.8	74.8
✓			✓	63.0	<u>65.8</u>	<u>67.4</u>	<u>69.5</u>	70.8	72.1	<u>74.6</u>
		✓	✓	<u>63.2</u>	65.0	66.9	69.1	71.2	71.3	74.3
		<i>p</i> -value		0.0324	0.0297	0.0331	0.0196	0.0203	0.0297	0.0406

Table 4: Investigation of different perturbation strategies in feature space. We report experimental results on DSB and DeepLesion using different perturbation strategies in feature space with various input resolutions. The first line corresponds to the results of CSD. The best results are in bold while the second best are underlined. **NRB** stands for the noisy residual block. Here we do not use adaptive consistency cost, cutout or adversarial perturbation. All p-values are calculated between the top-2 models in each column.

experiment for three times and report their average results.

4.5.1 An Incremental Study of Proposed Modules

In this part, we mainly conduct an incremental study of all proposed modules. Generally speaking, we can see that the proposed modules can consistently boost the overall performance in both DSB and DeepLesion datasets. More specifically, we can see that the adaptive consistency cost seems to bring larger improvements when compared to the noisy residual block and the instance-level perturbation, which are 1.5 points in DSB and 1.1 points in DeepLesion. After adding the noisy residual block, our SSMD can already outperform FocalMix by more than 0.5 point. Nonetheless, the instance-level adversarial perturbation finally helps SSMD to surpass FocalMix by a substantial margin in both DSB and DeepLesion datasets.

4.5.2 Investigation of Adaptive Consistency Cost

We study the influence of adding adaptive consistency cost in Table 3. It is obvious that the proposed adaptive consistency cost helps to boost the detection performance in all input resolutions. Specifically, the adaptive cost brings more than 1 point improvements on both DSB and DeepLe-

ξ	ϵ	Image Size				ξ	ϵ	Image Size		
		384	448	512	640			384	448	512
0	0	62.5	65.2	66.9	69.3	0	0	70.7	71.8	74.0
1e-8	1.0	63.0	65.6	67.3	69.7	1e-7	1.0	71.1	72.2	74.3
1e-7	1.0	63.2	65.8	67.6	70.2	1e-7	2.0	71.0	72.2	74.1
5e-7	1.0	<u>63.5</u>	<u>66.0</u>	<u>67.8</u>	<u>70.4</u>	5e-7	1.0	71.3	72.5	74.4
5e-7	2.0	63.7	66.2	68.2	70.5	5e-7	2.0	<u>71.5</u>	72.9	<u>74.7</u>
1e-6	2.0	63.4	65.9	67.4	70.1	1e-7	2.0	71.7	73.2	74.9
5e-7	3.0	63.3	65.8	67.2	70.3	1e-5	1.0	71.4	<u>73.0</u>	74.5
<i>p</i> -value		0.0743	0.0521	0.0634	0.0586	<i>p</i> -value		0.0627	0.0689	0.0754

(a) DSB
(b) DeepLesion

Table 5: Ablation study of ξ and ϵ in adversarial perturbation. When ξ or ϵ is 0, our method is equivalent to CSD [14]. We set τ to 0.95. Note that we do not use adaptive cost function, noisy residual block or cutout in these experiments. The best results are in bold while the second best are underlined. All *p*-values are calculated between the top-2 models in each column.

sion datasets. Interestingly, we can see that as the input resolution increases, the effectiveness of the adaptive cost becomes more significant. For example, on DSB dataset when the input resolution is 640×640 , the relative improvements are 1.8 points which are larger than those of small input resolutions. Similar phenomena also appears on DeepLesion dataset. The underlying reason can be that bigger inputs usually bring more region proposals where our adaptive mechanism is more effective in amplifying the significance of high-confidence proposals.

4.5.3 Investigation of Perturbation Strategies in Feature Space

In Table 4, we compare the performance of different perturbation strategies in feature space. Note that we did not use cutout, adaptive cost function or adversarial perturbation, so the results in Table 4 are lower than those in Table 1. The results of CSD is displayed in the first line of Table 4.

We first adopt the most widely used strategy Dropout [34] to regularize the feature space. Unfortunately, it seems that applying layer-wise dropout is not a good idea as it slightly degrades the

performance when compared to the original CSD. We then give a trial to Dropblock [9] which randomly masks rectangle regions in feature maps and was initially designed for object detection with natural images. Different from Dropout, Dropblock needs to be located at specific layers within the convolutional neural network. Following the instructions proposed by [9], we place Dropblock in conv4 and conv5 only. The results in Table 4 suggest that Dropblock achieves better performances than Dropout. Moreover, we try adding Gaussian noise (denoted as ‘Noise’) to each layer, which performs slightly better than Dropout but still worse than Dropblock.

The proposed layer-wise noisy residual block (denoted as ‘NRB’ in Table 4) outperforms Dropblock remarkably regardless of datasets and input resolutions. The reason behind may be that the strength of added noise in the noisy residual block is learnable. It means that different feature channels have different levels of noise. This principle enables the neural network to learn the optimal manner to implement layer-wise noise. In addition, we try to integrate Dropout or Gaussian noise with the proposed noisy residual block. However, we found that these naive combinations slightly degrade the detection performance.

4.5.4 Instance-level Adversarial Perturbation in Image Space

As described in Section 3.3, the proposed adversarial perturbation has three hyperparameters: the scale factor ξ , the confidence threshold τ and the magnitude controlling factor ϵ . In the following, we conduct comprehensive experiments to demonstrate the effect of each hyperparameter.

Scale Factor ξ and Confidence Threshold ϵ Table 5 displays how we determine the exact values of ξ and ϵ . In Table 5 the proposed method SSMD considerably surpasses the baseline CSD with a naive perturbation implementation. In practice, we determine the values of ξ and ϵ considering their influences on two datasets. For DSB dataset, the input resolution is set to 448×448 while the input size of DeepLesion is 512×512 . For other input sizes, we simply use the values of ξ and ϵ chosen in ablation study. As Table 5 displays, the setting of $\{\xi=5e-7, \epsilon=2.0\}$ achieves nearly the best performance in both datasets. Although the setting of $\{\xi=1e-7, \epsilon=2.0\}$ performs the best in DeepLesion dataset, we still adopt $\{\xi=5e-7, \epsilon=2.0\}$ as our default setting.

Confidence Threshold τ In Equation (9), we introduce a confidence threshold τ in order to control the influence of various proposals to generated adversarial examples. We show the effects of employing different confidence thresholds in Fig.5. It can be seen that as τ increases, using adversarial perturbation produces better performance. However, a too large value may lead to a

α	0.8	0.9	0.95	0.99	0.999	p -value
DSB	64.7	65.4	65.7	65.8	65.6	0.0347
DeepLesion	73.4	74.3	74.6	74.8	74.7	0.0286

Table 6: Investigation of EMA factor. The default input resolutions for DSB and DeepLesion are 448×448 and 512×512 , respectively. Note that we do not apply cutout or adversarial perturbation to network inputs and we also do not use adaptive cost function. All p -values are calculated between the top-2 models in each row.

c	x	y	w	h	DSB (mAP)				DeepLesion (4 FPs)		
					384	448	512	640	384	448	512
					60.0	62.3	65.0	67.8	67.7	69.2	71.5
✓					62.2	64.3	66.6	69.2	69.1	70.8	73.0
✓	✓				62.3	64.8	66.8	69.5	69.4	71.1	73.2
✓	✓	✓			62.9	65.6	67.2	69.9	70.6	72.0	74.3
✓	✓	✓	✓		62.9	65.7	67.4	69.9	71.0	72.4	74.5
✓	✓	✓	✓	✓	63.3	66.0	67.6	70.2	71.5	72.6	74.7
✓	✓	✓	✓	✓	62.5*	65.2*	66.9*	69.3*	70.7*	71.8*	74.0*
				p -value	0.0329	0.0272	0.0292	0.0314	0.0256	0.0236	0.0228

Table 7: Investigation of adding consistency loss with different combinations of predictions (c for confidence score, h for height, w for width, x and y for centering coordinate). * denotes the results of CSD. Note that we do not use cutout, noisy residual block and adversarial perturbation in these experiments. All p -values are calculated between the top-2 models in each column.

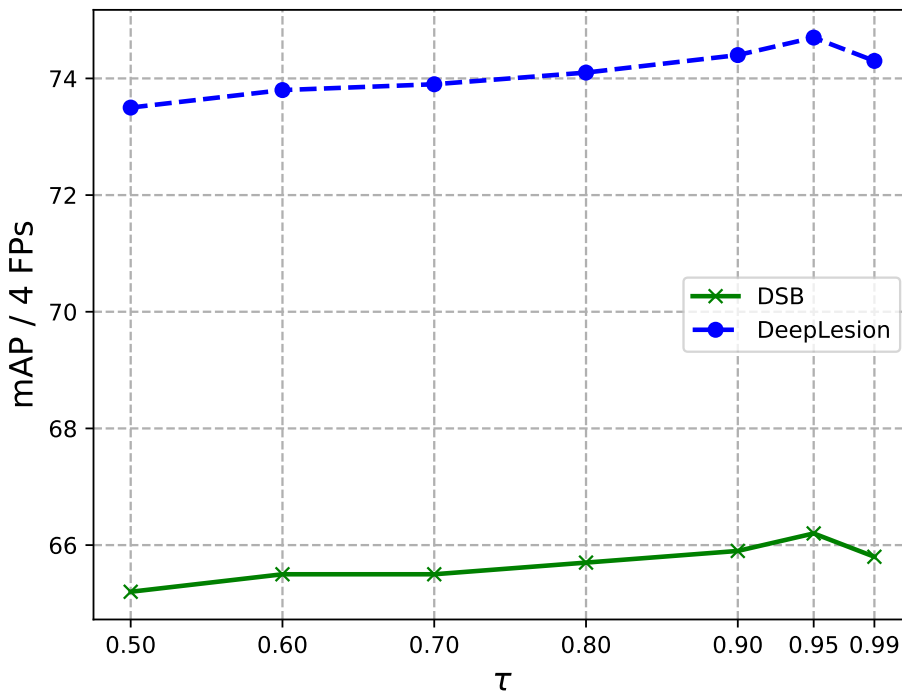


Figure 5: Investigation of different confidence thresholds. Here we do not apply cutout to network inputs. ξ is set to $5e-7$ while ϵ is set to 2.0. The default input resolutions for DSB and DeepLesion are 448×448 and 512×512 , respectively. The p -values between top-2 choices in DSB and DeepLesion are 0.0437 and 0.0348, respectively.

negative effect, such as 0.99. Such phenomenon is easy to explain since low-confidence regions might be incorrect and usually affect the detection accuracy negatively. In contrast, a threshold of 0.99 may filter out too many proposals. Based on the above observations, we set the default value of τ as 0.95 in both datasets.

4.5.5 Investigation of EMA Factor

Using EMA weights has been shown to help obtain more accurate predictions while stabilizing the training process [35, 44]. Compared to CSD, we replace the siamese network with a student-teacher architecture to promote the detection performance.

In Table 6, we gradually increase α in Equation (2) from 0.8 to 0.999 and report its corresponding performance in both DSB and DeepLesion datasets. Setting α as 0.99 seems to be optimal in both datasets. In contrast, decreasing α may deteriorate the detection accuracy. Particularly, a small

n	s	Image Size				n	s	Image Size		
		384	448	512	640			384	448	512
0	0	62.5	65.2	66.9	69.3	0	0	70.7	71.8	74.0
3	50	63.0	65.5	67.2	69.4	3	50	71.0	72.5	74.2
3	70	63.0	65.4	67.0	69.2	3	70	71.4	<u>72.8</u>	74.4
5	50	63.5	65.7	67.4	69.6	5	50	71.3	72.6	74.1
5	70	63.5	66.2	<u>67.7</u>	<u>69.8</u>	5	70	71.9	73.2	<u>74.5</u>
7	50	<u>63.3</u>	<u>65.8</u>	68.0	70.1	7	50	<u>71.6</u>	72.7	74.7
7	70	63.1	65.5	67.3	69.4	7	70	71.5	72.6	74.3
p -value		0.4893	0.1842	0.0765	0.0874	p -value		0.1445	0.0883	0.0981

(a) DSB

(b) DeepLesion

Table 8: Investigation of number (denoted as n) and size (denoted as s) of masks in cutout. We simply adopt rectangle masks of side length s in cutout. Note that in the above experiments, the baseline model ($n = 0$ and $s = 0$) corresponds to a vanilla CSD model. The best results are in bold while the second best are underlined. Note that we do not use adaptive cost function, noisy residual block or adversarial perturbation in these experiments. All p-values are calculated between the top-2 models in each column.

Dataset	1	3	5	p -value
DSB	64.3	67.1	68.8	0.0379
DeepLesion	72.8	75.0	76.5	0.0179

Table 9: Influence of adding consistency regularization to different numbers of feature scales. **1,3** and **5** denote the number of feature scales. The default input sizes of DSB and DeepLesion are 448×448 and 512×512 , respectively. All p-values are calculated between the top-2 models in each row.

value such as 0.8 may be harmful to the consistency based approach, leading to worse performance than a vanilla CSD (65.2 in DSB and 74.0 in DeepLesion). The underlying reason may be that the teacher network is updated too frequently which makes its predictions unstable to some degree. Oppositely, a too large value (0.999) may be harmful since the teacher network is updated too sparsely.

Dataset	SSD	Faster RCNN	Mask RCNN	RetinaNet (ours)	p -value
DSB	66.8	67.7	68.3	68.8	0.0438
DeepLesion	75.5	75.9	76.3	76.5	0.0379

Table 10: The default input resolutions for DSB and DeepLesion are 448×448 and 512×512 , respectively. All p -values are calculated between the top-2 models in each row. The labeled ratio is 50%.

Considering these observations, we set α to 0.99 in the rest of the experiments.

4.5.6 Investigation of Consistency Regularization

In Table 7, we present the performances of using different consistency regularization constraints. A simple baseline could be a model trained with normal and flipped images with manual annotations. Table 7 shows that adding consistency to predicted class scores improves the baseline. Specifically, simply regularizing the proposal confidence (denoted as c) can already bring more than 1 point improvement in both DSB and DeepLesion datasets. The largest improvements mainly come from the regularization of x and y . In most cases, regularizing these two variables outperforms regularizing the variable c only by approximate 1 point. The relative improvements in DeepLesion dataset are larger than those in DSB. Applying consistency regularization on all predicted variables brings about more than 3 points improvements over the supervised baseline. Compared with CSD, it can be observed that our proposed adaptive cost function leads to 1 point improvement.

4.5.7 Number and Size of Masks in Cutout

Cutout [7] is adopted in the proposed SSMD to add perturbations to input images. We apply cutout to the inputs of both student and teacher networks while adversarial perturbation is only applied to the input of teacher network.

As the number n and the size s (side length of a rectangle mask) of masks in cutout may affect the detection performance, it is necessary to conduct ablation studies for these two factors. In Table 8, the experimental results of different combinations of n and s are displayed. It is discovered that adding too many masks degrades the performance in both datasets. We argue that the reason can be




Figure 6: Visual comparison of different methods in DSB 2018. The colors of these approaches are consistent with those in Fig.3.




Figure 7: Visual comparison of different methods in DeepLesion. The colors of these approaches are consistent with those in Fig.3. Additionally, we use yellow boxes to denote the false positive results.

summarized as: more masks usually means hiding more instances which make the detector fail to predict the correct boxes. In Table 8 it is observed that as the input resolution increases, the effect of cutout becomes less significant. Specifically, $\{n=5, s=70\}$ and $\{n=7, s=50\}$ achieve two best and comparable results. $\{n=7, s=50\}$ seems to prefer larger input while $\{n=5, s=70\}$ affects more on




Figure 8: Proposal results from the student and the teacher model. For the left pair, we apply no consistency loss. For the right pair, we regularize the consistency of two models’ predictions. Note that for each image, we display 20 proposals with highest confidence scores before the non-maximum suppression (NMS) step. **w/o** and **w/** stand for without and with. These results come from the detectors trained for 50 epochs.

small input resolutions. We take $\{n=5,s=70\}$ as the default setting in most experiments.

4.5.8 Adding Consistency Loss to Different Numbers of Scales

To demonstrate the necessity of multi-scale consistency loss, we also report the experimental results of applying consistency loss to different numbers of scales in Table 9. It suggests that reducing the number of applied scales would adversely affect the overall detection performance. Specifically, applying consistency loss to 3 scales makes SSMD achieve comparable results with FocalMix. Interesting, even if we further reduce this number to 1, SSMD still performs much better than the supervised baseline, demonstrating the effectiveness of proposed modules.

4.5.9 Investigation of Different Detection Backbones

In Table 10, we investigate the influence of using different detection backbones. As Table 10 shows, Mask RCNN [12] achieves comparable results with our RetinaNet, although Mask RCNN runs much slower than our SSMD. Besides, Mask RCNN is better than Faster RCNN as it utilizes mask annotations during the training stage. SSD performs the worst across all 4 detection backbones, owing to its weakness in learning appropriate representations for medical image detection.

5 Visualization

In Fig.6 and Fig.7, we visually compare the supervised model, FocalMix, our SSMD and the fully-supervised baseline to the ground truth in DSB 2018 and DeepLesion, respectively. These results again verify the effectiveness of proposed SSMD. As Fig.6 and Fig.7 display, SSMD is superior to FocalMix in two different aspects: precision and recall. In DSB, SSMD is able to predict more accurate bounding boxes for nucleus. In contrast, FocalMix sometimes misses small nucleus while SSMD has the ability to localize these hard cases and hence achieves higher recall. Moreover, SSMD is stronger on detecting severely overlapped nucleus, which demonstrates the strength of its learned powerful image representations. When we turn to DeepLesion, it is obvious that SSMD again produces more accurate box predictions. Besides, SSMD also shows its strength in reducing false positive predictions. When comparing SSMD with the fully-supervised baseline, we can see that SSMD achieves acceptable results in most cases. Besides the detection results, we also display the proposal results of the student and the teacher model in Fig.8. The results indicate that the teacher models produce better detection results than those of student models in both pairs of images. If we compare their predictions carefully, it can be seen that the predictions of the student model with consistency loss are more accurate and consistent (with predictions of the teacher model) than those without consistency loss. These phenomena verify the findings from [1] which shows the teacher model helps to improve the learning process of the student model.

6 Conclusion

In this paper, we proposed a novel semi-supervised medical detection method which can boost the fully-supervised performance with additional unlabeled data. Specifically, the proposed detector consists of an adaptive consistency cost function, noisy residual blocks and an instance-level adversarial perturbation strategy. We also conduct experiments which not only demonstrate the strength of the overall proposed detector on various settings, but also verify the effectiveness of each single proposed module. In the future, we will explore more ways to improve the proposed method.

Credit Author Statement

Hong-Yu Zhou: Conceptualization; Investigation; Methodology; Software; Writing – original draft preparation & revision.

Chengdi Wang: Investigation; Visualization; Validation; Writing - original draft preparation.

Haofeng Li: Investigation; Writing – original draft preparation.

Gang Wang: Investigation; Resources.

Shu Zhang: Investigation; Validation.

Weimin Li: Supervision; Writing - review & editing.

Yizhou Yu: Conceptualization; Supervision; Writing - review & editing; Funding acquisition.

Acknowledgement

This work was funded in part by National Key Research and Development Program of China (No. 2019YFC0118101), National Natural Science Foundation of China (No. 91859203 and No. 82072005), Beijing Municipal Science and Technology Planning Project (No. Z211100003521009), Science and Technology Project of Chengdu (No. 2017-CY02–00030-GX), and Zhejiang Province Key Research & Development Program (No. 2020C03073). This research was conducted while Hong-Yu Zhou was visiting West China Hospital, Sichuan University.

Algorithm 1 Procedure of Semi-Supervised Medical Detection

Input:

- A batch of labeled images $\mathcal{X} = \{x^1, \dots, x^B\}$. A batch of unlabeled images $\tilde{\mathcal{X}} = \{\tilde{x}^1, \dots, \tilde{x}^B\}$.
 $\backslash\backslash B$ is the batch size
 $f_{\theta_s^n}(\cdot)$: Student network at time step n . $f_{\theta_t^n}(\cdot)$: Teacher network at time step n .
- 1: $\mathcal{X}_s = \text{cutout}(\text{Rot}(\mathcal{X}))$; $\tilde{\mathcal{X}}_s = \text{cutout}(\text{Rot}(\tilde{\mathcal{X}}))$ $\backslash\backslash$ Augmentation for the student network
 - 2: $\mathcal{X}_t = \text{Adv}(\text{cutout}(\text{Rot}(\text{Flip}(\mathcal{X}))))$; $\tilde{\mathcal{X}}_t = \text{Adv}(\text{cutout}(\text{Rot}(\text{Flip}(\tilde{\mathcal{X}}))))$ $\backslash\backslash$ Augmentation for the teacher network
 - 3: $\backslash\backslash$ Forward pass to get predictions
 - 4: **for** $n = 1$ to N **do** $\backslash\backslash N$ is number of training iterations
 - 5: $\theta_t^n = \alpha\theta_t^{n-1} + (1 - \alpha)\theta_s^n$ $\backslash\backslash$ Update teacher network
 - 6: $\text{loss}_{\text{cont}} = 0$; $\text{loss}_{\text{sup}} = 0$ $\backslash\backslash$ Initialize loss values
 - 7: **for** $i = 1$ to B **do**
 - 8: $p_s^c, p_s^x, p_s^y, p_s^w, p_s^h = f_{\theta_s^n}(x_s^i)$; $\tilde{p}_s^c, \tilde{p}_s^x, \tilde{p}_s^y, \tilde{p}_s^w, \tilde{p}_s^h = f_{\theta_s^n}(\tilde{x}_s^i)$ $\backslash\backslash$ Forward through student network
 - 9: $p_t^c, p_t^x, p_t^y, p_t^w, p_t^h = f_{\theta_t^n}(x_t^i)$; $\tilde{p}_t^c, \tilde{p}_t^x, \tilde{p}_t^y, \tilde{p}_t^w, \tilde{p}_t^h = f_{\theta_t^n}(\tilde{x}_t^i)$ $\backslash\backslash$ Forward through teacher network
 - 10: $\text{loss}_{\text{sup}} += \text{CE}(p_s^c, p_{gt}^c) + \text{SmoothL1}(p_s^{\{x,y,w,h\}}, p_{gt}^{\{x,y,w,h\}})$ $\backslash\backslash$ Apply supervised loss
 - 11: $\text{loss}_{\text{cont}} += \mathbf{W}(p_s^c, p_t^c) \otimes (\text{KL}(p_s^c, p_t^c) + \text{MSE}(p_s^{\{x,y,w,h\}}, p_t^{\{x,y,w,h\}}))$ $\backslash\backslash$ Consistency loss on labeled images
 - 12: $\text{loss}_{\text{cont}} += \mathbf{W}(\tilde{p}_s^c, \tilde{p}_t^c) \otimes (\text{KL}(\tilde{p}_s^c, \tilde{p}_t^c) + \text{MSE}(\tilde{p}_s^{\{x,y,w,h\}}, \tilde{p}_t^{\{x,y,w,h\}}))$ $\backslash\backslash$ Consistency loss on unlabeled images
 - 13: $\text{loss} = \text{loss}_{\text{sup}} + \lambda \text{loss}_{\text{cont}}$ $\backslash\backslash \lambda$ is a hyperparameter
 - 14: $\text{Backward}(\text{loss})$ $\backslash\backslash$ Update student network
-

References

- [1] Ben Athiwaratkun, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. There are Many Consistent Explanations of Unlabeled Data: Why You Should Average. *arXiv preprint arXiv:1806.05594*, 2018.
- [2] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *International Conference on Learning Representations*, 2019.
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. MixMatch: A Holistic Approach to Semi-Supervised Learning. In *Advances in Neural Information Processing Systems*, pages 5050–5060, 2019.
- [4] Qi Chang, Zhennan Yan, Yixuan Lou, Leon Axel, and Dimitris N Metaxas. Soft-Label Guided Semi-Supervised Learning for Bi-Ventricle Segmentation in Cardiac Cine MRI. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 1752–1755. IEEE, 2020.
- [5] Liang Chen, Paul Bentley, Kensaku Mori, Kazunari Misawa, Michitaka Fujiwara, and Daniel Rueckert. Self-supervised Learning for Medical Image Analysis using Image Context Restoration. *Medical image analysis*, 58:101539, 2019.
- [6] Shuai Chen, Gerda Bortsova, Antonio García-Uceda Juárez, Gijs van Tulder, and Marleen de Bruijne. Multi-Task Attention-Based Semi-Supervised Learning for Medical Image Segmentation. In *International Conference on Medical Image Computing and Computer Assisted Intervention*, pages 457–465. Springer, 2019.
- [7] Terrance DeVries and Graham W Taylor. Improved Regularization of Convolutional Neural Networks with Cutout. *arXiv preprint arXiv:1708.04552*, 2017.

- [8] Jia Ding, Aoxue Li, Zhiqiang Hu, and Liwei Wang. Accurate pulmonary nodule detection in computed tomography images using deep convolutional neural networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 559–567. Springer, 2017.
- [9] Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V Le. Dropblock: A Regularization Method for Convolutional Networks. In *Advances in Neural Information Processing Systems*, pages 10727–10737, 2018.
- [10] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [11] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and Harnessing Adversarial Examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2961–2969, 2017.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [14] Jisoo Jeong, Seungeui Lee, Jeessoo Kim, and Nojun Kwak. Consistency-based Semi-Supervised Learning for Object detection. In *Advances in Neural Information Processing Systems*, pages 10758–10767, 2019.
- [15] Zhanghan Ke, Daoye Wang, Qiong Yan, Jimmy Ren, and Rynson WH Lau. Dual Student: Breaking the Limits of the Teacher in Semi-Supervised Learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6728–6736, 2019.

- [16] Naji Khosravan and Ulas Bagci. S4nd: Single-shot single-scale lung nodule detection. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 794–802. Springer, 2018.
- [17] Samuli Laine and Timo Aila. Temporal Ensembling for Semi-Supervised Learning. *arXiv preprint arXiv:1610.02242*, 2016.
- [18] Bruno Lecouat, Ken Chang, Chuan-Sheng Foo, Balagopal Unnikrishnan, James M Brown, Houssam Zenati, Andrew Beers, Vijay Chandrasekhar, Jayashree Kalpathy-Cramer, and Pavitra Krishnaswamy. Semi-Supervised Deep Learning for Abnormality Classification in Retinal Images. *arXiv preprint arXiv:1812.07832*, 2018.
- [19] Dong-Hyun Lee. Pseudo-label: The Simple and Efficient Semi-Supervised Learning Method for Deep Neural Networks. In *ICML Workshop*, 2013.
- [20] Xiaomeng Li, Lequan Yu, Hao Chen, Chi-Wing Fu, and Pheng-Ann Heng. Semi-Supervised Skin Lesion Segmentation via Transformation Consistent Self-ensembling Model. *British Machine Vision Conference*, 2018.
- [21] Xiaomeng Li, Lequan Yu, Hao Chen, Chi-Wing Fu, Lei Xing, and Pheng-Ann Heng. Transformation-Consistent Self-Ensembling Model for Semisupervised Medical Image Segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [22] Zihao Li, Shu Zhang, Junge Zhang, Kaiqi Huang, Yizhou Wang, and Yizhou Yu. MVP-Net: Multi-view FPN with Position-aware Attention for Deep Universal Lesion Detection. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 13–21. Springer, 2019.
- [23] Fangzhou Liao, Ming Liang, Zhe Li, Xiaolin Hu, and Sen Song. Evaluate the malignancy of pulmonary nodules using the 3-d deep leaky noisy-or network. *IEEE Transactions on Neural*

Networks and Learning Systems, 30(11):3484–3495, 2019.

- [24] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature Pyramid Networks for Object Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2117–2125, 2017.
- [25] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal Loss for Dense Object Detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2980–2988, 2017.
- [26] Jingya Liu, Liangliang Cao, Oguz Akin, and Yingli Tian. 3dfpn-hs2: 3d feature pyramid network based high sensitivity and specificity pulmonary nodule detection. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 513–521. Springer, 2019.
- [27] Quande Liu, Lequan Yu, Luyang Luo, Qi Dou, and Pheng Ann Heng. Semi-Supervised Medical Image Classification with Relation-driven Self-ensembling Model. *IEEE Transactions on Medical Imaging*, 2020.
- [28] Ali Madani, Mehdi Moradi, Alexandros Karargyris, and Tanveer Syeda-Mahmood. Semi-Supervised Learning with Generative Adversarial Networks for Chest X-Ray Classification with Ability of Data Domain Adaptation. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 1038–1042. IEEE, 2018.
- [29] Ali Madani, Jia Rui Ong, Anshul Tibrewal, and Mohammad RK Mofrad. Deep Echocardiography: Data-efficient Supervised and Semi-Supervised Deep Learning towards Automated Diagnosis of Cardiac Disease. *NPJ digital medicine*, 1(1):1–11, 2018.
- [30] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual Adversarial Training: a Regularization Method for Supervised and Semi-Supervised Learning. *IEEE Transac-*

tions on Pattern Analysis and Machine Intelligence, 41(8):1979–1993, 2018.

- [31] Dong Nie, Yaozong Gao, Li Wang, and Dinggang Shen. ASDNet: Attention based Semi-Supervised Deep Networks for Medical Image Segmentation. In *International Conference on Medical Image Computing and Computer Assisted Intervention*, pages 370–378. Springer, 2018.
- [32] Aria Pezeshk, Sardar Hamidian, Nicholas Petrick, and Berkman Sahiner. 3-d convolutional neural networks for automatic detection of pulmonary nodules in chest ct. *IEEE Journal of Biomedical and Health Informatics*, 23(5):2080–2090, 2018.
- [33] Kihyuk Sohn, Zizhao Zhang, Chun-Liang Li, Han Zhang, Chen-Yu Lee, and Tomas Pfister. A simple semi-supervised learning framework for object detection. *arXiv preprint arXiv:2005.04757*, 2020.
- [34] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [35] Antti Tarvainen and Harri Valpola. Mean Teachers are Better Role Models: Weight-averaged Consistency Targets Improve Semi-Supervised Deep Learning Results. In *Advances in Neural Information Processing Systems*, pages 1195–1204, 2017.
- [36] Vikas Verma, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. Interpolation Consistency Training for Semi-Supervised Learning. *arXiv preprint arXiv:1903.03825*, 2019.
- [37] Dong Wang, Yuan Zhang, Kexin Zhang, and Liwei Wang. FocalMix: Semi-Supervised Learning for 3D Medical Image Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3951–3960, 2020.

- [38] Guo-Hua Wang and Jianxin Wu. Repetitive Reprediction Deep Decipher for Semi-Supervised Learning. *arXiv preprint arXiv:1908.04345*, 2019.
- [39] Ke Yan, Xiaosong Wang, Le Lu, and Ronald M Summers. DeepLesion: Automated Mining of Large-scale Lesion Annotations and Universal Lesion Detection with Deep Learning. *Journal of Medical Imaging*, 5(3):036501, 2018.
- [40] Xin Yi, Ekta Walia, and Paul Babyn. Unsupervised and Semi-Supervised Learning with Categorical Generative Adversarial networks Assisted by Wasserstein Distance for Dermoscopy Image Classification. *arXiv preprint arXiv:1804.03700*, 2018.
- [41] Lequan Yu, Shujun Wang, Xiaomeng Li, Chi-Wing Fu, and Pheng-Ann Heng. Uncertainty-aware Self-ensembling Model for Semi-Supervised 3D Left Atrium Segmentation. In *International Conference on Medical Image Computing and Computer Assisted Intervention*, pages 605–613. Springer, 2019.
- [42] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. Mixup: Beyond Empirical Risk Minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- [43] Yizhe Zhang, Lin Yang, Jianxu Chen, Maridel Fredericksen, David P Hughes, and Danny Z Chen. Deep Adversarial Networks for Biomedical Image Segmentation Utilizing Unannotated Images. In *International Conference on Medical Image Computing and Computer Assisted Intervention*, pages 408–416. Springer, 2017.
- [44] Hong-Yu Zhou, Shuang Yu, Cheng Bian, Yifan Hu, Kai Ma, and Yefeng Zheng. Comparing to Learn: Surpassing ImageNet Pretraining on Radiographs By Comparing Image Representations. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2020.

- [45] Yi Zhou, Xiaodong He, Lei Huang, Li Liu, Fan Zhu, Shanshan Cui, and Ling Shao. Collaborative Learning of Semi-Supervised Segmentation and Classification for Medical Images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2079–2088, 2019.
- [46] Zongwei Zhou, Vatsal Sodha, Md Mahfuzur Rahman Siddiquee, Ruibin Feng, Nima Tajbakhsh, Michael B Gotway, and Jianming Liang. Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 384–393. Springer, 2019.
- [47] Wentao Zhu, Chaochun Liu, Wei Fan, and Xiaohui Xie. Deeplung: Deep 3d dual path nets for automated pulmonary nodule detection and classification. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 673–681. IEEE, 2018.
- [48] Martin Zlocha, Qi Dou, and Ben Glocker. Improving Retinanet for CT Lesion Detection with Dense Masks from Weak Recist Labels. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 402–410. Springer, 2019.