

ElixirNet: Relation-Aware Network Architecture Adaptation for Medical Lesion Detection

Chenhan Jiang,^{2*} Shaoju Wang,^{1*} Hang Xu,² Xiaodan Liang,^{1†}
Nong Xiao¹

¹Sun Yat-Sen University, ²Huawei Noah's Ark Lab

Abstract

Most advances in medical lesion detection network are limited to subtle modification on the conventional detection network designed for natural images. However, there exists a vast domain gap between medical images and natural images where the medical image detection often suffers from several domain-specific challenges, such as high lesion/background similarity, dominant tiny lesions, and severe class imbalance. Is a hand-crafted detection network tailored for natural image undoubtedly good enough over a discrepant medical lesion domain? Is there more powerful operations, filters, and sub-networks that better fit the medical lesion detection problem to be discovered? In this paper, we introduce a novel ElixirNet that includes three components: 1) TruncatedRPN balances positive and negative data for false positive reduction; 2) Auto-lesion Block is automatically customized for medical images to incorporate relation-aware operations among region proposals, and leads to more suitable and efficient classification and localization. 3) Relation transfer module incorporates the semantic relationship and transfers the relevant contextual information with an interpretable graph, thus alleviates the problem of lack of annotations for all types of lesions. Experiments on DeepLesion and Kits19 prove the effectiveness of ElixirNet, achieving improvement of both sensitivity and precision over FPN with fewer parameters.

Introduction

Lesion detection in medical CT images is an important prerequisite for computer-aided detection/diagnosis (CADe/CADx). Recently, remarkable progress has been brought to the application of deep learning paradigms, especially Convolutional Neural Network (CNN) (Krizhevsky, Sutskever, and Hinton 2012), to CADe/CADx (Yan, Bagheri, and Summers 2018; Wang et al. 2017; Greenspan, Van Ginneken, and Summers 2016). Most of the works in this area directly use the natural image detection pipelines without adapting to the medical imaging domain. (Shin et al. 2016) transfers the CNN model with pre-trained ImageNet to medical images detection directly. (Jaeger et al.

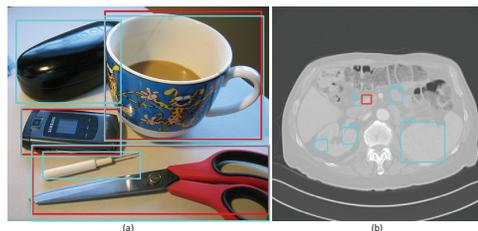


Figure 1: A vast domain gap between medical and natural images. FPN works well on MSCOCO (a), while it fails on DeepLesion (b) with the tiny-size lesion, identical nearby regions and overmuch false positive (ground-truth is red box).

2018) fuses RetinaNet detector (Lin et al. 2017b) with the U-Net (Ronneberger, Fischer, and Brox 2015) architecture to improve detection performance with full segmentation supervision. Moreover, (Yan, Bagheri, and Summers 2018) sends multiple neighbouring slices into the R-FCN (Dai et al. 2016) to generate feature maps separately, which are then aggregated for final prediction. Those methods finetune common detection framework on medical datasets without designing a special network for medical images. Different from natural images, only one to three lesions exist in a medical image and look similar to the nearby non-lesions. Also, lesions are usually tiny-sized and severe class imbalanced, illustrated in Figure 1. Using the conventional detection network for natural images directly is inefficient and will lead to a performance drop in those scenarios. Thus, customizing networks for medical lesion detection is in great need.

Conventional natural images detection pipelines most consist of three components: feature extractor; region proposal network (RPN) and RCNN head for processing proposals feature. While it is well-known that ImageNet pre-trained feature extractor (Russakovsky et al. 2015) is beneficial to medical images networks (Tajbakhsh et al.; Shin et al. 2016; 2016), RPN and RCNN head designed for natural images are inconsistent for medical images. To maintain a sufficiently high recall for proposals in natural images, uniformly sampling anchors over the whole localization in RPN leads to severe false candidates in medical im-

*Both authors contributed equally to this work.

†Corresponding Author: xliang328@gmail.com

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

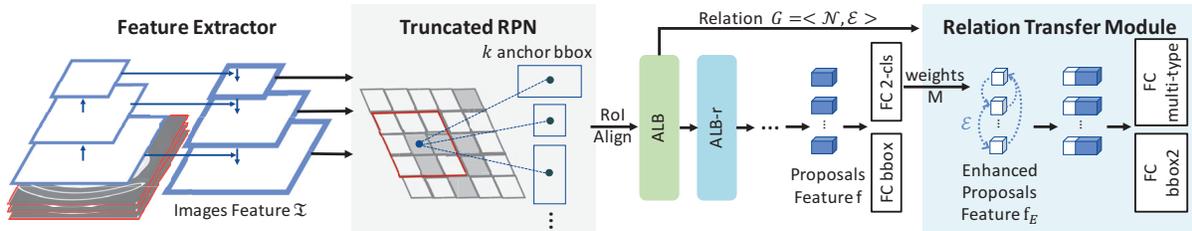


Figure 2: An overview of ElixirNet. TruncatedRPN is introduced to avoid overmuch false. Auto-lesion Blocks (ALB and ALB_r) are found by differentiable NAS to capture proposals representation and relation \mathcal{E} . Relation transfer module encodes more contextual information and transfers lesion presence embeddings M to enhanced proposal feature by learned relation \mathcal{E} .

ages. (Roth et al. 2015) trained classifiers on the aggregation of multiple 2D slices for false positive reduction, but it has two stages and is not end-to-end. We introduce TruncatedRPN (TRPN) to adaptively localize the lesion object for different input. RCNN heads in (Lin et al. 2017a; Ren et al. 2015) are too simple to exact discriminate lesion from the similar nearby non-lesion region. (Liu, Huang, and others 2018) considers various receptive fields to mimic human visual systems but its neglect of relational interaction. Relation Network (Hu et al. 2018) designed an adapted attention head to incorporate relationships from relevant proposals of which performance is limited by a single receptive field.

Recently neural architecture network search (NAS) has achieved much success in developing architectures that outperform hand-craft architectures on image classification and semantic segmentation task (Nekrasov et al.; de Moor et al.; Zoph et al. 2019; 2018; 2018). Different from (Zoph et al. 2018) on natural object detection that only transfers searched classification architecture as feature extractor and suffers from the large cost of GPU hours, we propose a novel search space for medical images following differentiable NAS (Liu, Simonyan, and Yang 2019), which optimizes architecture parameters based on the gradient descent. In this paper, we automatically customize an Auto-lesion block (ALB) as RCNN head substitute, which includes flexible receptive fields to locate the tiny-size lesion, relation context to identify similar regions and enough complexity to capture good feature representation.

Current lesion detection methods generally target one particular disease or lesion. Yet multiple types of lesions may appear on the same CT slice, e.g. abdomen/liver/kidney are neighboring types. Multiple correlated findings can help for better prediction and are crucial to building an automatic radiological diagnosis and reasoning system. However, it remains a challenge to develop a universal multi-purpose CAD framework due to rare annotations for all types of lesions, capable of detecting multiple disease types seamlessly. Compared with finetuning a new multi-type classifier, we introduce Relation Transfer Module (RTM) to incorporate the semantic relation learned by ALB and transfers the relevant contextual information with an interpretable graph in the end-to-end training. In a word, our ElixirNet consisted of TruncatedRPN, Auto-lesion block and relation transfer module can alleviate problems on medical lesion domain

and enable a unified lesion detection network seamlessly.

In our experiments on DeepLesion (Yan et al. 2018a), we observe consistent gains on the 3 different common feature extractor with reducing almost one-third of the parameters. The search is very efficient even for 512×512 input and only takes about 2 days on 4 GTX1080. The sensitivity of lesion detection with 3 false positives per image has about 2.3% improvement based on FPN. And for multi-type lesion detection, consistent improvement of sensitivity on all kinds of lesion can be found.

Related Work

Automated Lesion Detection. The detection of lesions in images is a key part of diagnosis and is one of the most labor-intensive for clinicians. Typically, the tasks consist of the localization and identification of small lesions in the full image space. Different from common object detection, lesion detection needs to consider 3D context information while facing more challenging problems about serve imbalance of categories and tiny-size objects. Recently, great progress has been brought to medical imaging analysis by deep learning methods (Liao et al.; Yan et al.; Yu et al. 2019; 2018b; 2018). Many conventional methods which are based on hand-crafted local features have been replaced by deep neural networks, which typically produces higher detection accuracy (Greenspan, Van Ginneken, and Summers 2016). However, most lesion detection (Jaeger et al.; Wang et al.; Yan, Bagheri, and Summers; Yan et al. 2018; 2017; 2018; 2018a) enhanced by CNN directly extend existing detection frameworks, which are typically designed for natural images, to process volumetric CT data. (Liao et al. 2019) uses small 3D patches of CT data as the input of the region proposal network (RPN) (Ren et al. 2015). Another method (Yan, Bagheri, and Summers 2018) sends multiple neighboring slices into R-FCN (Dai et al. 2016) to generate feature maps separately, which are then aggregated for final prediction. These methods do not focus on designing specific network for medical lesion detection.

Neural Architecture Search Method. Neural Architecture Search (NAS) aims at automatically designing finding the optimal architecture in a given search space such that the validation accuracy is maximized on the given task. Conventional architecture search methods can be categorized as random with weights prediction (Brock et al. 2017), evo-

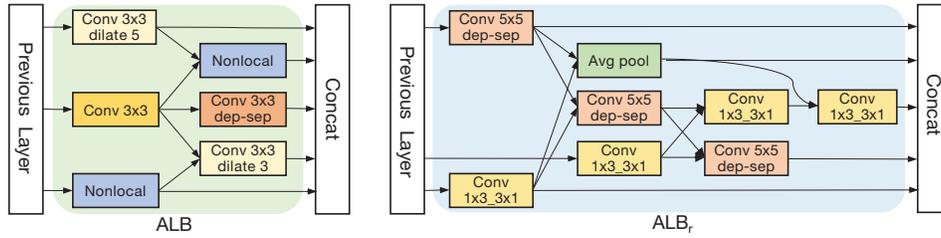


Figure 3: The detailed architecture of ALB and ALB_r . ALB keeps the same spatial resolution and output channels with input including flexible dilated convolution to locate tiny-size lesion and Non-local operator to incorporate relation-aware information among lesions. To deepen network, an inception-like block ALB_r is designed for double channels and resolution reduction.

lution (Miikkulainen et al. 2019), and reinforcement learning (Zoph and Le 2017; Zoph et al. 2018), which have achieved highly competitive performance on image classification. Only few of them focus on object detection. DetNAS (Chen et al. 2019) use the one-shot NAS to design a new feature extractor, while ImageNet pretrained feature extractor is beneficial to medical image networks. NAS-FPN (Ghiasi, Lin, and Le 2019) only focuses on feature fusion connection in the feature extractor and uses RL-based search algorithm. The search is very time-consuming even though a proxy task with ResNet-10 backbone. More efficient solutions for NAS have been recently proposed. Instead of searching over a discrete set of candidate architectures, (Liu, Simonyan, and Yang 2019; Liu et al. 2019) relax the search space to be continuous, so that the architecture can be optimized concerning its validation set performance by gradient descent. Our work follows the differentiable NAS formulation and extends it into the specific medical lesion detection setting.

The Proposed Approach

Our ElixirNet is based on two stage detection pipeline with feature pyramid extractor. Since lesion detection relies on 3D context information and current detectors are mostly designed with a three channels input (RGB), we naturally group the neighboring three slices to 3-channel images and thus implement current detection backbone with pretrained weights on ImageNet (Russakovsky et al. 2015). As shown in Figure 2, we introduce TruncatedRPN to avoid overmuch false positive prediction by restraining some regions where lesions are unlikely to exist. Then the output of RoIAlign is feed into Auto-lesion block (ALB), which is found by differentiable NAS according to a novel search space based on observation of many popular designs and medical images characteristics. After aggregating information from various receptive fields and learning a region-to-region undirected graph $G : G = \langle \mathcal{N}, \mathcal{E} \rangle$ in ALB, we further propose a relation transfer module to encode more contextual information and transfer coarse embeddings of lesion existing to enhance proposal feature by learned relation \mathcal{E} . In what follows, we describe the structure of the above mentioned three modules and explain their implementation in detail.

TruncatedRPN

RPN is basis in two-stage nature image object detection pipelines, which takes image feature from feature extractor

and outputs a set of rectangular object proposals with objectness score by sliding-window at each location. And it supposes a uniform distribution of anchor localization for all input. This method is inefficient for medical images, as rare lesion objects exist on the image and many of anchors are placed in regions where lesion objects are impossible to arise. In this work, we aims to develop a more efficient and suitable TruncatedRPN (TRPN) for lesion detection which can adaptively localize the lesion object for different input.

Denote \mathcal{I} as the images feature extracted from the input image, (i, j) as specific location in \mathcal{I} , original RPN can be formulate as $\mathcal{F}(\mathcal{I}(i, j)|k)$, k is anchor prior related to location (i, j) like scale and aspect ratio settings. Following this formulation, TRPN introduces a module to predict distribution of anchor localization $p(\mathcal{I}(i, j))$ and truncates the localition whose predicted confidence are below predefined threshold ϵ , as follows:

$$TRPN = \mathcal{F}([p(\mathcal{I}(i, j)) \cdot \mathcal{I}(i, j)]_{p(\mathcal{I}(i, j)) > \epsilon} | k) \quad (1)$$

We empirically apply 1×1 convolution with an element-wise sigmoid function as p . And we set mean of predicted distribution as ϵ . This mechanism can filter out half of the regions and avoid redundant information.

Auto-lesion Block

The proposed Auto-lesion Block (ALB) is a medical data-friendly block with hybrid receptive fields and kernel size. The Figure3 shows the detailed architecture of ALB with keeping the same spatial resolution and output channels with input, and ALB_r with reducing the half of spatial resolution and increasing double channels. In ALB, dilated convolution with flexible receptive fields can locate the tiny-size lesion, and Non-local operator captures relation-aware information to distinguish lesions with similar no-lesion regions. And ALB_r is a inception-like block, mentioned in (Szegedy et al. 2016), e.g. 3×3 convolution is composed in a asymmetric 1×3 and 3×1 convolution. It also uses large kernel depth-wise convolution and average pooling to keep a trade-off between computational overhead and network complexity. These two blocks are customized by medical data-oriented NAS with the same search space and strategy.

Search Space for ALB Assume that ALB is a directed acyclic graph consisting of \mathcal{B} branches, each branch has

Method	Sensitivity IoU (%)							mAP (%)	Speed (img/ms)	#_param (M)
	0.5	1	2	3	4	6	8			
3DCE, 3 slices (Yan, Bagheri, and Summers 2018)	56.49	67.65	76.89	80.56	82.76	85.71	87.03	25.6	42	17.2
3DCE, 9 slices (Yan, Bagheri, and Summers 2018)	59.32	70.68	79.09	82.78	84.34	86.77	87.81	27.9	56	19.2
3DCE, 27 slices (Yan, Bagheri, and Summers 2018)	62.48	73.37	80.70	83.61	85.65	87.26	89.09	29.2	114	25.2
FRCNN, 3 slices (Ren et al. 2015)	55.94	66.25	74.69	79.17	79.87	79.87	79.87	27.9	54	33.0
FPN, 3 slices (Lin et al. 2017a)	60.57	70.34	77.65	81.51	83.55	85.67	85.67	31.5	68	41.4
Deformv2, 3 slices (Zhu et al. 2019)	63.16	72.35	81.04	84.47	86.52	87.76	87.76	34.3	74	42.2
ElixirNet, 3 slices	60.43 ^{+4.49}	69.38 ^{+3.13}	77.20 ^{+2.51}	80.52 ^{+1.35}	82.72 ^{+2.85}	83.74 ^{+3.9}	83.74 ^{+3.9}	30.8 ^{+2.9}	78	19.9
ElixirNet w FPN, 3 slices	66.39 ^{+5.82}	75.41 ^{+5.07}	82.05 ^{+4.4}	86.63 ^{+4.12}	87.75 ^{+4.2}	88.36 ^{+2.69}	88.36 ^{+2.69}	35.2 ^{+3.7}	90	28.4
ElixirNet w Deformv2, 3slices	65.64 ^{+2.48}	75.84 ^{+3.49}	83.70 ^{+2.66}	86.92 ^{+2.45}	88.91 ^{+2.39}	91.37 ^{+3.61}	91.37 ^{+3.61}	36.3 ^{+2.0}	102	29.3

Table 1: Comparison with state-of-the-art object detection on DeepLesion. Sensitivity at various FPs per image on the test set of the official data split of DeepLesion. IoU criteria and 3 slices input (one key slice and two neighbor slices) are used.

2 inputs from previous branches and outputs 1 tensor following (Liu, Simonyan, and Yang 2019), and initial channels of each branch is D . Branch b can be specified as a 5-tuple $(X_1, X_2, OP_1, OP_2, Y_b)$, where $X_1, X_2 \in \mathcal{X}_b$ specific input tensors, $OP_1, OP_2 \in \mathcal{OP}$ specific operations to apply to the corresponding inputs. The output Y concatenates with all branches outputs $Y = \text{concat}(Y_1, Y_2, \dots, Y_b)$ with $B \times D$ channels. In most of the image classification NAS frameworks (Zoph and Le 2017; Zoph et al. 2018; Liu, Simonyan, and Yang 2019), similar candidate operations for \mathcal{OP} are considered. These pre-defined operations sets may be reasonable in the natural image classification task, while medical lesion object detection depends on semantic relation among proposals and flexible receptive fields. In addition, (Liu, Huang, and others 2018) show the relationship between the size and eccentricity of receptive fields can enhance the discriminability and robustness of the feature representation. Inspired by the above methods, we consider 9 operators as the set of candidate operations in \mathcal{OP} :

- no connection
- skip connection
- 3×3 average pooling
- Non-local
- 1×3 and 3×1 depthwise-separable conv
- 3×3 depthwise-separable conv
- 5×5 depthwise-separable conv
- 3×3 atrous conv w dilate rate 3
- 3×3 atrous conv w dilate rate 5

All operations are of stride 1 and the convolved feature maps are padded to preserve their spatial resolution. We use the ReLU-Conv-BN order for convolutional operations. Our search space contains 3×10^{-9} structures.

Non-local Operator. The Non-local operator aims to encode semantic relation between region proposals which is relevant to the object detection. We formulate the relation as a region-to-region undirected graph $\mathcal{G} : \mathcal{G} = \langle \mathcal{N}, \mathcal{E} \rangle$, where each node in \mathcal{N} corresponds to a region proposal and each edge $e_{i,j} \in \mathcal{E}$ encodes relationship between two nodes. Formally, the input of non-local operator is the proposal features $X \in \mathbb{R}^{N \times D}$ from previous branches. The adjacency matrix for the undirected graph \mathcal{G} with self-loops can be then calculated by a matrix multiplication as:

$$\mathcal{E} = \text{softmax}(\phi(X)\phi(X)^T), \quad (2)$$

where $\phi(\cdot)$ is non-linear transformation with ReLU activation. Then we use a layer GCN (Kipf and Welling 2017) for the propagation of each node in the graph \mathcal{E} . A simple form is $Y = \sigma(\mathcal{E}\psi(X)W)$, where W and $\psi(\cdot)$ are non-linear transformation, and σ is active function. In this paper,

we consider a fully-connected layer with ReLU activation as the non-linear function and output dimension is $\frac{D}{2}$. For the output of this operation, we use another fully-connected layer g to keep the shape of the input tensor. The final output of the Non-local operator is $g(Y)$.

Differentiable Search Strategy The search strategy of ALB builds on a continuous relaxation of the discrete architectures described in (Liu, Simonyan, and Yang 2019). The branch’s output tensor Y_i is a weighted mixture of the operations with $|\mathcal{OP}|$ parallel paths which are connected to all elements in X_i :

$$Y_i = \sum_{Y_j \in X_i} \sum_{OP_k \in \mathcal{OP}} \alpha_{ij}^k OP_k(Y_j) \quad (3)$$

where the weight $\{\alpha_{ij}^k\}$ are the architecture parameters calculated by applying softmax to $|\mathcal{OP}|$, k is in $\{ALB, ALB_r\}$. After the continuous relaxation, the task of architecture search reduces to learning a set of continuous variables $\{\alpha_{ij}^k\}$. Therefore they can be optimized efficiently using stochastic gradient descent (SGD). We randomly split training data into two disjoint sets, and the loss function includes cross-entropy about classification and smooth $L1$ loss about location, calculated on these sets are denoted by \mathcal{L}_{train} and \mathcal{L}_{val} . The optimization alternates between:

1. Update network weights ω by $\nabla_{\omega} \mathcal{L}_{train}(\omega, \alpha)$
2. Update architecture weights α by $\nabla_{\alpha} \mathcal{L}_{val}(\omega, \alpha)$

At the end of search, a discrete architecture is obtained by replacing each mixed operation OP_{ij} with the most likely operation by argmax: $OP_{ij} = \text{argmax}_{OP_k \in \mathcal{OP}} \alpha_{ij}^k$.

Relation Transfer Module

To develop a universal multi-purpose CAD framework, we introduce Relation Transfer Module (RTM) to endow multi-type lesion detection. Rather than simply finetuning a new multi-type classifier, we transfer coarse embeddings of lesion existing to fine lesion types information and try to propagate semantic embeddings among proposals by learned relation \mathcal{E} in ALB.

In some zero/few-shot problems, (Salakhutdinov, Torralba, and Tenenbaum 2011; Gong et al. 2018) use the classifier’s weights as the high-level semantic embedding or representation of category. Thus the collection of the weights $M \in \mathbb{R}^{2 \times (P+1)}$ of original binary classifier (including the

Method	Sensitivity IoBB (%)						
	0.5	1	2	3	4	6	8
3DCE, 3 slices (Yan, Bagheri, and Summers 2018)	58.43	70.95	80.64	85.12	87.30	90.72	92.37
3DCE, 9 slices (Yan, Bagheri, and Summers 2018)	62.64	74.37	83.29	86.99	89.15	91.61	92.90
3DCE, 27 slices (Yan, Bagheri, and Summers 2018)	64.01	75.69	83.71	87.52	88.25	91.47	92.85
FRCNN, 3 slices (Ren et al. 2015)	57.59	68.46	77.36	81.96	82.43	82.43	82.43
FPN, 3 slices (Lin et al. 2017a)	62.17	72.35	80.29	83.82	86.03	87.68	87.68
Deform, 3 slices (Zhu et al. 2019)	65.09	74.49	83.34	86.84	88.86	89.71	89.71
ElixirNet, 3 slices	63.36 ^{+5.77}	73.21 ^{+4.75}	81.38 ^{+4.02}	85.06 ^{+3.10}	86.53 ^{+4.10}	87.50 ^{+5.07}	87.50 ^{+5.07}
ElixirNet w FPN, 3 slices	68.67 ^{+6.50}	77.93 ^{+5.58}	84.83 ^{+4.54}	88.62 ^{+4.80}	90.59 ^{+4.56}	91.77 ^{+4.09}	91.77 ^{+4.09}
ElixirNet w Deformv2, 3slices	67.33 ^{+2.24}	77.89 ^{+3.40}	86.00 ^{+2.66}	90.13 ^{+3.29}	91.66 ^{+2.80}	93.55 ^{+3.84}	93.55 ^{+3.84}

Table 2: Sensitivity with IoBB overlap criteria at various FPs per image on the test set of the official data split of DeepLesion.

Method	Sensitivity IoU (%)			mAP (%)
	0.5	1	2	
FPN (Lin et al. 2017a)	53.68	56.71	58.46	21.1
Deformv2 (Zhu et al. 2019)	67.18	68.31	71.21	32.9
ElixirNet w FPN	54.69 ^{+1.01}	57.93 ^{+1.22}	61.20 ^{+2.74}	29.6 ^{+8.5}
ElixirNet w Deformv2	67.28 ^{+0.10}	69.02 ^{+0.71}	80.05 ^{+2.4}	33.6 ^{+0.7}

Table 3: Comparison on Kits19 with 3 slices input.

%	Lesion Type							
	BN	AB	ME	LV	LU	KD	ST	PV
FPN w finetune (Lin et al. 2017a)	-	79.58	71.34	78.05	83.32	-	57.96	72.21
ElixirNet w/o RTM	56.52	84.14	83.41	82.99	87.10	60.70	75.67	80.64
ElixirNet	57.38	85.41	84.62	83.94	88.72	65.51	77.24	83.07

Table 4: Multi-type lesion detection of sensitivity at 3 FPs per image on the test set. The abbreviations of lesion types stand for bone, abdomen, mediastinum, liver, lung, kidney, soft tissue, pelvis and bone, respectively. IoU as overlap computation criteria is used. “FPN w finetune” and “ElixirNet w/o RTM” finetune new multi-type classifier on validation datasets. “-” means too low recall to evaluate sensitivity.

bias) is regarded as category-wise semantic embeddings, P is the output dimension of the ALB_r . Since our graph G is a region-to-region graph extracted from Non-local operator in ALB, we need to find most appropriate mappings from category-wise semantic embeddings to region-wise representations of node $f_i \in \mathbf{f}$ (the input of RTM). For avoiding bias produced by original binary classification, we use a soft-mapping which compute the mapping weights $\gamma_{M \rightarrow x_i} \in \Gamma^s$ as $\gamma_{M \rightarrow x_i} = \frac{\exp(s_{ij})}{\sum_j \exp(s_{ij})}$, where s_{ij} is the classification score for the region i towards category j from the previous binary classification layer, denoted “FC 2-cl” in Figure 2. Then the process of graph reasoning can be solved by matrix multiplication: $f_E = \mathcal{E}\Gamma^s M W_E$, where $W_E \in \mathbb{R}^{N \times \mathbb{E}}$ is a transformation weight matrix and E is the output dimension of the RTM module. Finally, the enhanced feature f_E is concatenated to the original region features \mathbf{f} to improve both classification and localization of the multi-type lesion.

Experiments

Datasets and Evaluations. We conduct experiments on the DeepLesion (Yan et al. 2018a) and Kits19 (Heller et al. 2019) datasets. DeepLesion is a large-scale dataset on

876,934 axial CT slices (mostly 512×512) from 10,594 CT studies of 4,427 unique patients. DeepLesion contains a variety of lesions (Yan et al. 2018b). However, only 32,120 axial slices have bounding boxes annotation. There are 1–3 lesions in each axial slice, totally 32,735 lesions. We select 30% as validation (4889 lesions) and test (4927 lesions), while the rest is regarded as the training set (22919 images) following the official division (Yan, Bagheri, and Summers 2018; Yan et al. 2018a). The lesions in validation and test sets are all labeled with specific types, and have been categorized into the 8 subtypes of lung (2394 lesions), abdomen (2176 lesions), mediastinum (1672 lesions), liver (1284 lesions), pelvis (681 lesions), soft tissue (867 lesions), kidney (495 lesions), and bone (247 lesions). The lesions in the training set are annotated by bounding box but without labels of lesion types. To validate the generalization capability of the model, we conduct experiments on Kits19 (Heller et al. 2019) which is a kidney tumor semantic segmentation challenge. We random split 80% patients (totally 210 kidney cancer patients) as our training set (13147 images) and the left as testing set (3209 images). Since Kits19 is a segmentation dataset, we convert segmentation masks to bounding boxes for all lesion.

For all the evaluation, we adopt mean Average Precision (mAP) across IoU 0.5-0.9 as evaluation, which is designed for natural images detection following (Lin et al. 2017a). Sensitivity (IoU ≥ 0.5) with different false positives per image is a commonly used metric in medical detection (Yan, Bagheri, and Summers 2018). Considering the prediction may still be viewed as a true positive (TP), it can also help the radiologists when IoU is less than 0.5. To overcome this evaluation bias, we also utilized the intersection over the detected bounding-box area ratio (IoBB) as another criterion, following (Yan, Bagheri, and Summers 2018).

Implementation Details. We conduct all experiments on a single server with 4 GTX1080 cards in Pytorch (Paszke et al. 2017). Three widely-adopted natural images detection methods i.e. FRCNN (Ren et al. 2015), FPN (Lin et al. 2017a) and Deformv2 (Zhu et al. 2019) are regarded as baseline network to show the generalization ability of our ElixirNet. Unless otherwise noted, settings are the same for all experiments.

TruncatedRPN. Based on the original RPN, we empirically apply 1×1 convolution with an element-wise sigmoid function on RPN feature map to obtaining the confidence

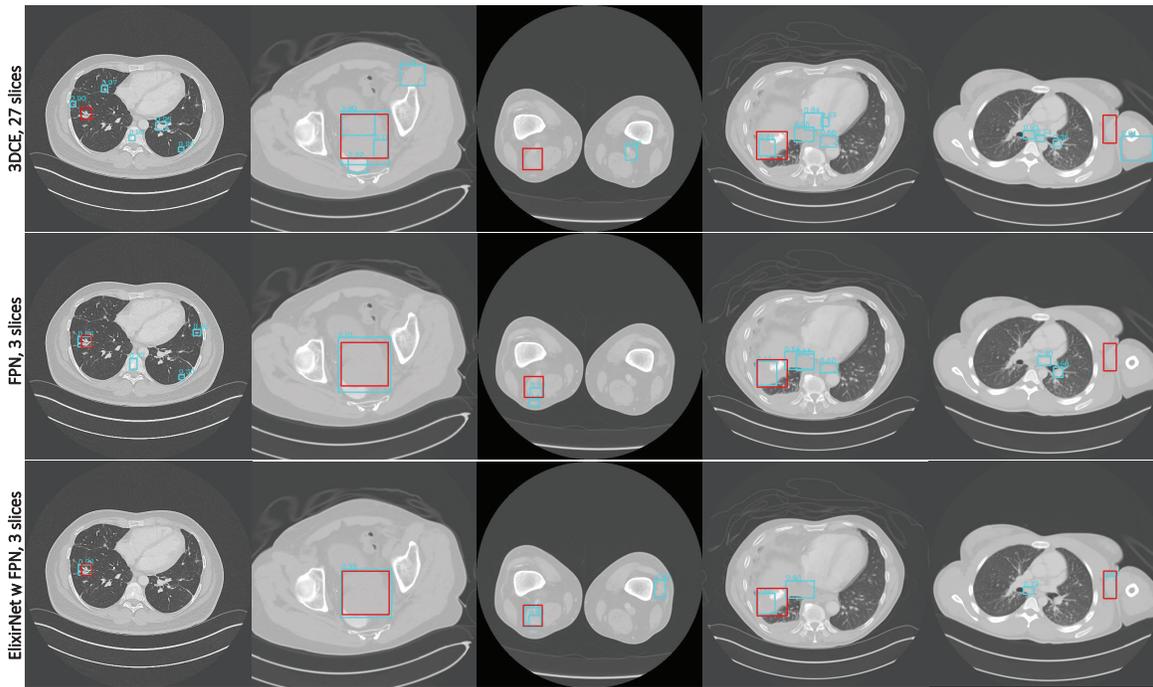


Figure 4: Qualitative results on DeepLesion among 3DCE with 27 slices, FPN and our ElixirNet with FPN. Detection results with higher confidence, precise location and less false positive can be achieved by our method. The prediction with confidence score > 0.3 is visualized. The location of ground-truth (red box) is shown for easy identification.

map of the locations. Region proposals are then generated at the locations where the confidence is greater than the mean of the confidence map. Five anchor scales (2, 3, 4, 6, 12) and three aspect ratios of (0.5, 1, 2) are used in TRPN following (Yan, Bagheri, and Summers 2018).

Auto-Lesion Block Search. Owing to memory restriction, feature extractor and RPN are frozen with initializing pretrained baseline network on DeepLesion. We consider the branch $\mathcal{B} = 4$ and the initial channels D of each branch is 16. To carry out architecture search, we split half of the DeepLesion training data as the *val* set. Network weights ω is updated after training 15 epochs with batch size = 16 (same with validation sets). We choose momentum SGD with initial learning rate 0.02 (annealed down to zero following a cosine schedule), momentum 0.9, and weight decay 3×10^{-4} as an optimizer for weights ω , while architecture weights α are optimized by Adam with initial learning rate 3×10^{-4} , momentum 0.999 and weight decay 10^{-3} .

ElixirNet. Series connection of feature extractor, TRPN, 1 ALB, 2 ALB_r and RTM forms our final ElixirNet. The detailed architecture of ALB and ALB_r can be found in Figure 3 and we double initial channels D in Auto-Lesion Block Search as channels of each operator, and use a 3×3 convolution to reduce output feature channels to 128 before feeding to ALB. And we set $P = 512$, $E = 128$ in RTM. For network training, stochastic gradient descent (SGD) is performed on 4 GPUs per 2 images for 12 epochs. ResNet-50 (He et al. 2016) pretrained on ImageNet (Russakovsky et al. 2015) is used. The initial learning rate is 0.02, reduce three

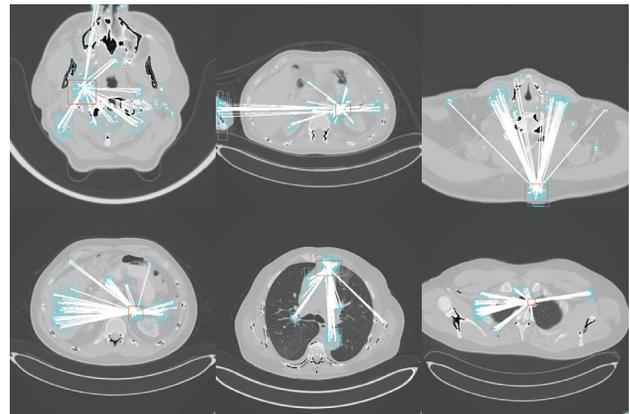


Figure 5: Visualization of learned relation G from ALB. The centers of regions are plotted and connected by the learned edges \mathcal{E} . Edges thickness correspond to the strength of the graph edge weights. Note other regions are linked to regions which are close to ground-truth (red box) location.

times ($\times 0.01$); 10^{-4} as weight decay; 0.9 as momentum.

Comparison with state-of-the-art

Single Lesion Detection Benchmarks. The results of Lesion Detection are presented in Table 1 and Table 2. We compare with 3DCE (Yan, Bagheri, and Summers 2018), FRCNN (Ren et al. 2015), FPN (Lin et al. 2017a) and

%	TRPN	ALB	RTM	input 9 slices	3 FPs (IoU) sensitivity	mAP
FPN (Lin et al. 2017a)					81.51	31.5
	✓				83.61	33.9
		✓			83.88	34.1
			✓		84.19	34.0
	✓	✓			84.72	34.8
	✓		✓		85.43	34.4
ElixirNet	✓	✓	✓		86.63	35.2

Table 5: Ablation Studies on sensitivity and mAP. The impact of different modules and input slices is explored.

Deformv2 (Zhu et al. 2019) with FPN backbone. Notably, ElixirNet achieved significant improvements than all natural images baselines with fewer parameters. Our ElixirNet improves around 10.3% for FRCNN, 11.7% for FPN and 5.8% for Deformv2 on mAP respectively. And the result shows ElixirNet works mostly on small false positive per images (improving 5.82% of IoU overlap 6.50% of IoBB overlap sensitivity at 0.5 FPs on FPN). Our method with high accuracy at low error rates is consistent with the needs of radiologists. The parameter size of ElixirNet is about 60%-70% of the baseline detection network. Without optimization of parallel structure in Pytorch, inference speed is slightly slower than baseline. We also evaluate Kits19(Heller et al. 2019) to compare with FPN (Lin et al. 2017a) and Deformv2 (Zhu et al. 2019) with FPN backbone. The results can be found in Table 3. As can be seen, our method outperforms 40.3% and 2% than FPN and Deformv2 on mAP, and has to continue improvements on sensitivity with IoU. Figure 4 shows quality results comparison of the 3DCE with 27 slices input, FPN and our ElixirNet with FPN on DeepLesion. Our method is more accurate than the 3DCE and baseline FPN due to relation-aware architecture adaptation for medical images. The graph structure learned from ALB is in Figure 5. The proposed ALB learned interpretable relation among regions. Other regions are linked to regions which are close to the ground-truth location. The relation-aware knowledge helps improve proposals feature thus leading to better detection performance. **Multi-type Lesion Detection.** Different lesion types of labelling are provided only in the test and validation set of DeepLesion (Yan et al. 2018a). A simple solution is replacing binary classifier and bounding-box regression with new multi-type classifier and bounding-box regression, denoted as “FPN w finetune” and “ElixirNet w/o RTM”. During the training phase, feature extractor and RPN are frozen, while RCNN head and our ALB finetune for 6 epochs with 0.002 learning rate on the validation set. The results in Table 4 show that our “ElixirNet w/o RTM” far outperforms FPN on generalization and robustness of multi-type lesion detection. Furthermore, the results with RTM achieve an average 2% sensitivity of all types with only increasing 0.2% parameters.

Ablative Analysis. To analyze the importance of different modules in Table 5, ablation studies are conducted on FPN (Lin et al. 2017a) baseline. All results are on sensitivity with IoU overlap criteria at 3 FPs per image and mAP. 1) ALB is the most vital component with relation-aware adaptation

Method and Modifications	sensitivity (%)	mAP (%)
ElixirNet	83.88	34.1
ALB w/o our search space (Liu, Simonyan, and Yang 2019)	81.06 ^{-2.82}	32.1 ^{-2.0}
random ALB (Li and Talwalkar 2019)	81.84 ^{-2.04}	32.1 ^{-2.0}

Table 6: Comparison of different search space and strategy. The input with 3 slices and evaluation on sensitivity at 3 FPs (IoU criteria) and mAP are used for all methods. “w/o our search space” means using the search space in image classification NAS works directly. “random ALB” replaces random search strategy with our differentiable search strategy.

for medical images. 2) TRPN can improve performance for false positive reduction and outperforms focal loss (Lin et al. 2017b) with replacing the cross-entropy loss in RPN whose sensitivity at 3 FPs (with IoU criteria) is 82.66%. 3) The usage of RTM further improves the sensitivity by 1.91% and mAP by 0.4%. 4) More slices are considered following (Yan, Bagheri, and Summers 2018); the performance of the 9 slices is slightly higher than the 3 slices. Since 72% of lesions can be covered by 3 slices input in DeepLesion dataset, our final model uses 3 slices input. **Search space and strategy in ALB.** The comparison from Table 6 shows that the proposed novel search space is most significant for designing specific architecture for medical images. Original search space in previous NAS work causes 2.82% sensitivity falloff and 2% decrease on mAP. Also, differentiable search strategy searching architectures based on the gradient descent is a benefit for medical lesion detection.

Conclusion

In this work, we proposed a novel ElixirNet which is customized for medical lesion detection with a composition of TruncatedRPN, Auto-lesion Block and Relation transfer module. It can adaptively suppress anchor location with interested lesion absence and captures the semantic context of a key proposal from relation-aware neighbourhoods, leading to more suitable and efficient prediction and false positive reduction. The stable and consistent performance of our ElixirNet on all evaluation criteria of DeepLesion and Kits19 outperforms current methods with fewer parameters.

Acknowledgments. This research is partly supported by the National Natural Science Foundation of China (No. U1611461) and in part by the National Natural Science Foundation of China (NSFC) under Grant No.61976233.

References

- Brock, A.; Lim, T.; Ritchie, J. M.; and Weston, N. 2017. Smash: one-shot model architecture search through hypernetworks. *arXiv preprint arXiv:1708.05344*.
- Chen, Y.; Yang, T.; Zhang, X.; Meng, G.; Pan, C.; and Sun, J. 2019. Detnas: Neural architecture search on object detection. In *NeurIPS*.
- Dai, J.; Li, Y.; He, K.; and Sun, J. 2016. R-fcn: Object detection via region-based fully convolutional networks. In *NerIPS*.
- de Moor, T.; Rodriguez-Ruiz, A.; Mérida, A. G.; Mann, R.; and Teuwen, J. 2018. Automated soft tissue lesion detection and segmentation in digital mammography using a u-net deep learning network. In *IWBI*.

- Ghiasi, G.; Lin, T.-Y.; and Le, Q. V. 2019. Nas-fpn: Learning scalable feature pyramid architecture for object detection. In *CVPR*.
- Gong, C.; He, D.; Tan, X.; Qin, T.; Wang, L.; and Liu, T.-Y. 2018. Frage: Frequency-agnostic word representation. In *NerIPS*.
- Greenspan, H.; Van Ginneken, B.; and Summers, R. M. 2016. Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique. *IEEE Transactions on Medical Imaging* 35(5):1153–1159.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Heller, N.; Sathianathan, N.; Kalapara, A.; Walczak, E.; Moore, K.; Kaluzniak, H.; Rosenberg, J.; Blake, P.; Rengel, Z.; Oestreich, M.; et al. 2019. The kits19 challenge data: 300 kidney tumor cases with clinical context, ct semantic segmentations, and surgical outcomes. *arXiv preprint arXiv:1904.00445*.
- Hu, H.; Gu, J.; Zhang, Z.; Dai, J.; and Wei, Y. 2018. Relation networks for object detection. In *CVPR*.
- Jaeger, P. F.; Kohl, S. A.; Bickelhaupt, S.; Isensee, F.; Kuder, T. A.; Schlemmer, H.-P.; and Maier-Hein, K. H. 2018. Retina u-net: Embarrassingly simple exploitation of segmentation supervision for medical object detection. *arXiv preprint arXiv:1811.08661*.
- Kipf, T. N., and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *NerIPS*.
- Li, L., and Talwalkar, A. 2019. Random search and reproducibility for neural architecture search. *arXiv preprint arXiv:1902.07638*.
- Liao, F.; Liang, M.; Li, Z.; Hu, X.; and Song, S. 2019. Evaluate the malignancy of pulmonary nodules using the 3-d deep leaky noisy-or network. *IEEE transactions on neural networks and learning systems*.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017a. Feature pyramid networks for object detection. In *CVPR*.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017b. Focal loss for dense object detection. In *ICCV*.
- Liu, C.; Chen, L.-C.; Schroff, F.; Adam, H.; Hua, W.; Yuille, A.; and Fei-Fei, L. 2019. Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation. In *CVPR*.
- Liu, S.; Huang, D.; et al. 2018. Receptive field block net for accurate and fast object detection. In *ECCV*.
- Liu, H.; Simonyan, K.; and Yang, Y. 2019. Darts: Differentiable architecture search. In *ICLR*.
- Miikkulainen, R.; Liang, J.; Meyerson, E.; Rawal, A.; Fink, D.; Francon, O.; Raju, B.; Shahrzad, H.; Navruzyan, A.; Duffy, N.; et al. 2019. Evolving deep neural networks. In *Artificial Intelligence in the Age of Neural Networks and Brain Computing*. Elsevier.
- Nekrasov, V.; Chen, H.; Shen, C.; and Reid, I. 2019. Fast neural architecture search of compact semantic segmentation models via auxiliary cells. In *CVPR*.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch. In *NerIPS Workshop*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NerIPS*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*.
- Roth, H. R.; Lu, L.; Liu, J.; Yao, J.; Seff, A.; Cherry, K.; Kim, L.; and Summers, R. M. 2015. Improving computer-aided detection using convolutional neural networks and random view aggregation. *IEEE transactions on medical imaging* 35(5):1170–1181.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115(3):211–252.
- Salakhutdinov, R.; Torralba, A.; and Tenenbaum, J. 2011. Learning to share visual appearance for multiclass object detection. In *CVPR*.
- Shin, H.-C.; Roth, H. R.; Gao, M.; Lu, L.; Xu, Z.; Nogues, I.; Yao, J.; Mollura, D.; and Summers, R. M. 2016. Deep convolutional neural networks for computer-aided detection: Cnn architectures, dataset characteristics and transfer learning. *IEEE transactions on medical imaging* 35(5):1285–1298.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *CVPR*.
- Tajbakhsh, N.; Shin, J. Y.; Gurudu, S. R.; Hurst, R. T.; Kendall, C. B.; Gotway, M. B.; and Liang, J. 2016. Convolutional neural networks for medical image analysis: Full training or fine tuning? *IEEE transactions on medical imaging* 35(5):1299–1312.
- Wang, Z.; Yin, Y.; Shi, J.; Fang, W.; Li, H.; and Wang, X. 2017. Zoom-in-net: Deep mining lesions for diabetic retinopathy detection. In *MICCAI*.
- Yan, K.; Bagheri, M.; and Summers, R. M. 2018. 3d context enhanced region-based convolutional neural network for end-to-end lesion detection. In *MICCAI*.
- Yan, K.; Wang, X.; Lu, L.; and Summers, R. M. 2018a. Deeplesion: automated mining of large-scale lesion annotations and universal lesion detection with deep learning. *Journal of Medical Imaging* 5(3):036501.
- Yan, K.; Wang, X.; Lu, L.; Zhang, L.; Harrison, A. P.; Bagheri, M.; and Summers, R. M. 2018b. Deep lesion graphs in the wild: relationship learning and organization of significant radiology image findings in a diverse large-scale lesion database. In *CVPR*.
- Yu, Q.; Xie, L.; Wang, Y.; Zhou, Y.; Fishman, E. K.; and Yuille, A. L. 2018. Recurrent saliency transformation network: Incorporating multi-stage visual cues for small organ segmentation. In *CVPR*.
- Zhu, X.; Hu, H.; Lin, S.; and Dai, J. 2019. Deformable convnets v2: More deformable, better results. In *CVPR*.
- Zoph, B., and Le, Q. V. 2017. Neural architecture search with reinforcement learning. In *ICLR*.
- Zoph, B.; Vasudevan, V.; Shlens, J.; and Le, Q. V. 2018. Learning transferable architectures for scalable image recognition. In *CVPR*.