

Efficient Classification of Very Large Images with Tiny Objects

Fanjie Kong, Ricardo Henao

Department of Electrical and Computer Engineering
 Duke University

{fanjie.kong, ricardo.henao}@duke.edu

Abstract

An increasing number of applications in computer vision, specially, in medical imaging and remote sensing, become challenging when the goal is to classify very large images with tiny informative objects. Specifically, these classification tasks face two key challenges: *i*) the size of the input image is usually in the order of mega- or giga-pixels, however, existing deep architectures do not easily operate on such big images due to memory constraints, consequently, we seek a memory-efficient method to process these images; and *ii*) only a very small fraction of the input images are informative of the label of interest, resulting in low region of interest (ROI) to image ratio. However, most of the current convolutional neural networks (CNNs) are designed for image classification datasets that have relatively large ROIs and small image sizes (sub-megapixel). Existing approaches have addressed these two challenges in isolation. We present an end-to-end CNN model termed Zoom-In network that leverages hierarchical attention sampling for classification of large images with tiny objects using a single GPU. We evaluate our method on four large-image histopathology, road-scene and satellite imaging datasets, and one gigapixel pathology dataset. Experimental results show that our model achieves higher accuracy than existing methods while requiring less memory resources.

1. Introduction

Neural networks have achieved state-of-the-art performance in many image classification tasks [24]. However, there are still many scenarios where neural networks can still be improved. Using modern deep neural networks on image inputs of very high resolution is a non-trivial problem due to the challenges of scaling model architectures [44]. Such images are common for instance in satellite or medical imaging. Moreover, these images tend to become even bigger due to the rapid growth in computational and memory availability, as well as the advancements in camera sensor technology. Specifically challenging are the so called

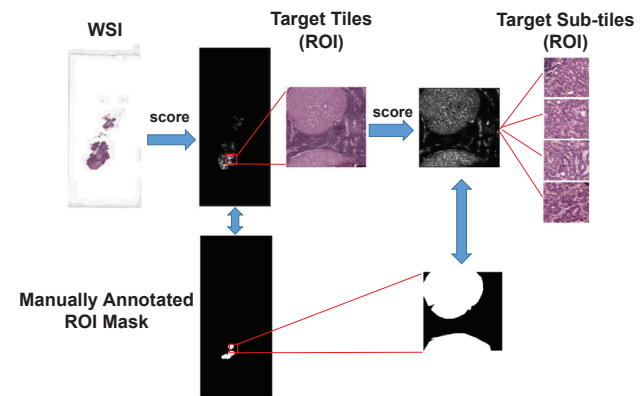


Figure 1. Illustration of processing a typical WSI using our zoom-in strategy. We see that *i*) there are large regions with little information (mostly background), and *ii*) small informative regions have high-resolution details. Leveraging the above characteristics of WSI, we derive a method that gradually zooms-in to the ROI. The proposed approach first process the down-sampled WSI to sample the target tiles, and then repeats this procedure to sample target sub-tiles. The sampled sub-tiles contain the fine-grained information for classification. The bottom images show that the manually annotated ROIs are captured by the proposed approach without the need for pixel level annotations.

tiny object image classification tasks, where the goal is to classify images based on the information of very small objects or regions of interest (ROIs), in the presence of a much larger and rich (*non-trivial*) background that is uncorrelated or non-informative of the label. Consequently, constituting an input image with a very low ROI-to-image ratio. Recent work [37] showed that with a dataset of limited size, convolutional neural networks (CNNs) have poor performance on very low ROI-to-image ratio problems. In these settings, the input resolution is increased from typical image sizes, *e.g.*, 224×224 pixels, to *gigapixel* images of size ranging from $45,056 \times 35,840$ to $217,088 \times 111,104$ pixels [30], which not only require significantly more computational processing power per image than a typical image given a fixed deep architecture, but in some cases, become prohibitive for current GPU-memory standards.

Figure 1 shows an example of a gigapixel pathology image, from which we see that manually annotated ROIs (with cancer metastases), not usually available for model training, constitute a small proportion of the whole slide image (WSI). Moreover, many tasks in satellite imagery [5] and medical image analysis [30] are still challenging due to the scarce methodology available for such big images.

Other recent works have addressed the computational resource bottlenecks associated with models for very large images by proposing approaches such as the streaming neural network [38] and gradient checkpoint [33]. However, these methods do not take advantage of the characteristics of very large images in tiny object image classification tasks, *i.e.*, those in which only a small portion of the image input is informative for the classification label of interest. Alternatively, other approaches use visual attention models to exploit these characteristics and show that discriminative information may be sparse and scattered across various image scales [12, 20, 36], which suggests that in some scenarios, processing the entire input image is unnecessary, and specially true in tiny object image classification tasks. For instance, [20] leverages attention to build image classifiers using a small collection of tiles (image patches) *sampled* from the matrix of attention weights generated by an attention network. Unfortunately, despite the ongoing efforts, existing approaches are either prohibitive or require severe resolution trade-offs that ultimately affect classification performance, for tasks involving very large (gigapixel) images.

The purpose of this work is to address these limitations simultaneously. Specifically, we propose a neural network architecture termed *Zoom-In network*, which as we will show, yields outperforming memory efficiency and classification accuracy on various tiny object image classification datasets. We build upon [20] by proposing a two-stage *hierarchical* attention sampling approach that is effectively able to process gigapixel images, while also leveraging contrastive learning as a means to improve the quality of the attention mechanisms used for sampling. This is achieved by building aggregated representations over a small fraction of high-resolution content (sub-tiles) that is selected from an attention mechanism, which itself leverages a lower resolution view of the original image. In this way, the model can dramatically reduce the data acquisition and storage requirements in real-world deployments. This is possible because low resolution views can be used to indicate which regions of the image should be acquired (attended) at higher resolution for classification purposes, without the need of acquiring the entire image at full resolution. Moreover, we show that the proposed approach can be easily extended to incorporate pixel-level-annotations when available for additional performance gains. Results on five challenging datasets demonstrate the capabilities of the Zoom-In network in terms of accuracy and memory efficiency.

2. Zoom-In Network

Below we present the construction of the proposed Zoom-In network model, which aims to efficiently process gigapixel images for classification of very large images with tiny objects. We start by briefly describing the one-stage attention sampling method proposed in [20], which we leverage in our formulation. Then, we introduce our strategy consisting in decomposing the attention-based sampling into two stages as illustrated in Figure 2. This *two-stage hierarchical sampling* approach enables computational efficiency without the need to sacrifice performance due to loss of resolution, when used in applications with very large images and small ROI-to-image ratios. In the experiments, we will show that the Zoom-In network results in improved performance relative to existing approaches on several tiny object image classification datasets, and importantly, without the need for any pixel-level annotation.

2.1. Attention Sampling

Let $T_{s_1}(x, c)$ denote a function that extracts a tile of size $h_1 \times w_1$ from the input, full-resolution, image $x \in \mathbb{R}^{H \times W}$ corresponding to the location (coordinates) $c = \{i, j\}$ in a lower resolution view $V(x, s_1) \in \mathbb{R}^{h \times w}$ of x at scale $s_1 \in (0, 1)$, so $h = \lfloor s_1 H \rfloor$ and $w = \lfloor s_1 W \rfloor$, where $\lfloor \cdot \rfloor$ is the floor operator. More specifically, $T_{s_1}(x, c)$ maps c to a location in x via $\{\lfloor 1 + (i - 1)(W - 1)/(w - 1) \rfloor, \lfloor 1 + (j - 1)(H - 1)/(h - 1) \rfloor\}$, and returns a tile of size $h_1 \times w_1$. Note that *i)* the map of locations between $V(x, s_1)$ and x only depends on the size of x ($H \times W$) and s_1 and not on the tile size ($h_1 \times w_1$); *ii)* $h_1, w_1 > 1/s_1$, to guarantee full coverage of x ; *iii)* this strategy requires to zero-pad x on all sides by $\lfloor h_1/2 \rfloor$ and $\lfloor w_1/2 \rfloor$ pixels accordingly; and *iv)* we have omitted the (color) channel dimension in x and $V(x, s_1)$ for notational simplicity, however, we consider color images (with an additional dimension) in our experiments. Then, let $\Psi_{\Theta}(x) = g_{\Theta}(f_{\Theta}(T_{s_1}(x, c)))$ be a neural network parameterized by Θ whose intermediate representation $z \in \mathbb{R}^K$ is obtained via feature extracting function $z = f_{\Theta}(T_{s_1}(x, c))$, *e.g.*, a convolutional neural network (CNN). Further, $g_{\Theta}(z)$ is a classification function also specified as a neural network and parameterized by Θ . We can provide $\Psi_{\Theta}(x)$ with an attention mechanism as follows

$$\alpha = a_{\Theta}(V(x, s_1)) : \mathbb{R}^{h \times w} \rightarrow \mathbb{R}^{h \times w} \quad (1)$$

$$\Psi_{\Theta}(x) = g_{\Theta} \left(\sum_{c \in C} \alpha_c f_{\Theta}(T_{s_1}(x, c)) \right), \quad (2)$$

where α is the matrix of attention weights such that $\sum_{c \in C} \alpha_c = 1$, $a_{\Theta}(V(x, s_1))$ is the attention function, also specified as a neural network, and C (of length $|C| = h \cdot w$) is the collection of all index pairs for view $V(x, s_1)$. In order to avoid computing the features z from all $|C|$ tiles implied by view $V(x, s_1)$, which can be a very large number




Figure 2. Illustration of the Zoom-In network. In Stage I, attention network $a_{\Theta}(\cdot)$ generates an attention map for the input image down-scaled by s_1 , from which N tiles are sampled with replacement (see samples map). In Stage II, attention network $b_{\Theta}(\cdot)$ generates an attention map for each selected tile and selects a sub-tile, thus N sub-tiles are selected (without replacement). Then all sub-tiles are fed to feature extractor $f_{\Theta}(\cdot)$, feature maps are aggregated using their corresponding attention weights, and predictions are obtained from aggregated features using a classification module $g_{\Theta}(\cdot)$. Further, both attention maps are also used to draw contrastive samples with minimal computational overhead (during training).

if x is big, as in our pathology and remote sensing scenarios, [20] proposed to leverage Monte Carlo estimation by only considering a small set of tiles from the original input image sampled via the attention function. This strategy leverages that α defines a discrete distribution over the set of $|C|$ tiles. Specifically, [20] approximates (2) by sampling from (1) via

$$\Psi_{\Theta}(x) \approx g_{\Theta} \left(\frac{1}{N} \sum_{c \in Q} f_{\Theta}(T_{s_1}(x, c)) \right), \quad (3)$$

where Q is a collection of $N \ll |C|$ index pairs for view $V(x, s_1)$ drawn independently and identically distributed (iid) from the distribution defined by the attention weights, i.e., $Q = \{(i, j) \sim a_{\Theta}(V(x, s_1)) | i = 1, 2, \dots, N\}$. In [20], they consider tiles of size $h_1 = w_1 = 27$, $s_1 = 0.2$ and $N = 10$ for the colon cancer dataset. See Experiments below for additional details.

Using the approximation in (3), the attention mechanism uses a lower resolution view $V(x, s_1)$ of the original image x for computing the attention distribution and outputs an aggregated feature vector by averaging over the features $\{z_n\}_{n=1}^N$ of a small amount of N tiles. Unfortunately, this approach is still prohibitive for gigapixel images because feasible combinations of h_1 , w_1 and s result in unrealistic memory needs for current GPU-memory standards. Below, we introduce the proposed two-stage hierarchical sampling to improve the memory efficiency of attention sampling.

2.2. Two-stage Hierarchical Attention Sampling

Multistage and hierarchical sampling strategies are often preferred in practice. For instance, the cost of interviewing or testing people are enormously reduced if these people are geographically or organizationally grouped, thus sampling is performed within groups (clusters). Such sampling design has many real-world applications such as household and mortality surveys, as well as high-resolution remote sensing applications [6, 13, 50]. Motivated by this idea, we design a two-stage hierarchical sampling approach to

reduce memory requirements when processing very large, gigapixel, images without severe resolution trade-offs.

Specifically, let $V(x, s_2, c) \in \mathbb{R}^{u \times v}$ be a view of $T_{s_1}(x, c)$ at scale $s_2 \in (0, 1)$, so $u = \lfloor s_2 h \rfloor = \lfloor s_1 s_2 H \rfloor$ and $v = \lfloor s_2 w \rfloor = \lfloor s_1 s_2 W \rfloor$. Further, we define a function $T_{s_2}(T_{s_1}(x, c), c')$ that extracts a sub-tile of size $h_2 \times w_2$ at location $c' = \{i', j'\}$ in $V(x, s_2)$ from tile $T_{s_1}(x, c)$ at location $c = \{i, j\}$ in $V(x, s_1)$. The mapping function $T_{s_2}(T_{s_1}(x, c), c')$ is defined similarly to $T_{s_1}(x, c)$, but returns tiles of size $h_2 \times w_2$ instead of $h_1 \times w_1$, and is such that $h_2 < h_1$, $w_2 < w_1$, and $h_2, s_2 > 1/s_2$. Moreover, we can also define an attention mechanism for $V(x, s_2, c)$ as in (1) as follows

$$\beta^c = b_{\Theta}(V(x, s_2, c)) : \mathbb{R}^{u \times v} \rightarrow \mathbb{R}_+^{u \times v}, \quad (4)$$

where β is the matrix of attention weights for the tile at location c of $V(x, s_1)$ such that $\sum_{c' \in C'} \beta_{c'}^c = 1$, $b_{\Theta}(V(x, s_2, c))$ is the attention function, also specified as a neural network, and C' (of length $|C'| = u \cdot v$) is the collection of all index pairs for view $V(x, s_2, c)$ of $T_{s_1}(x, c)$. Provided that $\sum_{c \in C} \alpha_c = 1$ in (1) and $\sum_{c' \in C'} \beta_{c'}^c = 1$ in (4), it is easy to see that $\sum_{c \in C} \sum_{c' \in C'} \alpha_c \beta_{c'}^c = 1$ and that the attention for location $c' = \{i', j'\}$ in $V(x, s_2, c)$ relative to the entire image x is $\alpha_c \beta_{c'}^c$. Consequently, we can rewrite (2) as

$$\Psi_{\Theta}(x) = g_{\Theta} \left(\sum_{c \in C} \alpha_c \sum_{c' \in C'} \beta_{c'}^c f_{\Theta}(T_{s_2}(T_{s_1}(x, c), c')) \right), \quad (5)$$

where now, the aggregated representation is a weighted average of all tiles of size $h_2 \times w_2$ of x , and like in (3), we can approximate as

$$\Psi_{\Theta}(x) \approx g_{\Theta} \left(\frac{1}{N} \sum_{c \in Q} f_{\Theta}(T_{s_2}(T_{s_1}(x, c), c')) \right), \quad (6)$$

where $c' \sim b_{\Theta}(V(x, s_2, c))$ is drawn iid from distribution $b_{\Theta}(V(x, s_2, c))$ for every location $c \in Q$.

Note that the approximation in (6) uses *full-resolution* sub-tiles from x that are drawn hierarchically from the two-level discrete distribution implied by α and $\{\beta^c\}_{c=1}^{|C|}$, which are obtained from *low-resolution* views $V(x, s_1)$ and $\{V(x, s_2, c)\}_{c=1}^{|C|}$. Importantly, in practice we do not need to instantiate the tiles $T_{s_1}(x, c)$ but only $T_{s_2}(T_{s_1}(x, c), c')$, and the second-level attention matrix in (4) can be obtained as needed (*on the fly*). However, this can cause computational inefficiency if multiple samples from the same location c are selected for level-two sampling in (6). Inefficiency occurs because such procedure will require to instantiate view $V(x, s_2, c)$ multiple times to obtain β^c on a single model update (iterations), and then when a sub-tile is sampled multiple times when obtaining $f_{\Theta}(T_{s_2}(T_{s_1}(x, c), c'))$. We can mitigate the inefficiency by ordering the samples in Q to prevent recalculating β^c , and we can reuse features $f_{\Theta}(T_{s_2}(T_{s_1}(x, c), c'))$ for a given c and c' as needed. Alternatively, we can avoid reusing sub-tiles by sampling locations c' in (6) *without replacement*. However, such sampling strategy will not be iid and as a result, it will cause bias in the Monte Carlo approximation in (6). Fortunately, using a formulation similar to that of [20], we can still obtain an unbiased estimator of the average (expectation) in (6) from a non-iid sample, without replacement, by leveraging the Gumbel-Top- k trick [23], which is extended from the Gumbel-Max trick for weighted reservoir sampling [11]. Specifically, from (5) we can write

$$\mathbb{E}_{c' \sim b_{\Theta}(V(x, s_2, c))}[f_{\Theta}(T_{s_2}(T_{s_1}(x, c), c'))] = \sum_{c' \in C'} \beta_{c'}^c f_{\Theta}(T_{s_2}(T_{s_1}(x, c), c')), \quad (7)$$

from which we can see that the sum on the right is an unbiased estimator of the expectation on the left. Alternatively, we can write

$$\mathbb{E}_{c' \sim b_{\Theta}(V(x, s_2, c))}[f_{\Theta}(T_{s_2}(T_{s_1}(x, c), c'))] = \sum_{c' \in C'} \sum_{i \neq c'} \beta_{c'}^c \frac{\beta_i^c}{1 - \beta_{c'}^c} (\beta_{c'}^c f_{\Theta}(T_{s_2}(T_{s_1}(x, c), c'))) + (1 - \beta_{c'}^c) f_{\Theta}(T_{s_2}(T_{s_1}(x, c), i)), \quad (8)$$

where $\beta_i^c / (1 - \beta_{c'}^c)$ is the attention weight for the i -th sub-tile reweighted to exclude sub-tile c' , which is equivalent to having already sampled it. The proof of (8) can be found in the Supplementary Material (SM). We can then approximate (5) like in (6) but sampling without replacement using

$$\Psi_{\Theta}(x) \approx g_{\Theta} \left(\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{i-1} \beta_{c_j'}^{c_i} f_{\Theta}(T_{s_2}(T_{s_1}(x, c_i), c_j')) \right) + \left[1 - \sum_{j=1}^{i-1} \beta_{c_j'}^{c_i} \right] f_{\Theta}(T_{s_2}(T_{s_1}(x, c_i), c_j')), \quad (9)$$

and c_j' is sampled via

$$c_j' \sim p(c' | c'_1, \dots, c'_{j-1}) \propto \begin{cases} \beta_i^c & \text{if } i \notin \{c'_1, \dots, c'_{j-1}\} \\ 0 & \text{otherwise} \end{cases},$$

where $p(c' | c'_1, \dots, c'_{j-1})$ represents sampling location c_j' without replacement, by having already sampled locations c'_1, \dots, c'_{j-1} .

Memory requirements In practice, the memory requirements of the attention sampling model are determined by the model parameters, feature maps, gradient maps and workspace variables [41]. For neural-network-based image models, memory allocation is mainly dominated by the size of the input image, *i.e.*, H and W . Specifically, the *peak memory* usage at inference for N samples scales with $\mathcal{O}(s^2 HW + N h_2 w_2)$ and $\mathcal{O}(s_1^2 HW + N' s_1^2 s_2^2 HW + N h_2 w_2)$ for both, the one-stage [20] and the proposed two-stage hierarchical model. Here, we use N' to denote the number of unique tiles in Q and s to indicate the scale of the view for the one-stage approach. In fact, we can show that our model requires significantly less GPU memory than one-stage attention sampling by choosing $s_1 < s$ and $s_2 = s$. Note that the number of selected tiles in the first stage decreases dramatically as the attention map is being optimized. We use the term *peak memory* to refer to the worse case scenario. Empirically, we have observed that the average number of selected tiles is $N' \approx N/2$. A detailed analysis of memory requirements is presented in the SM.

2.3. Efficient Contrastive Learning with Attention Sampling

Motivated by [22], we introduce a contrastive learning objective for the proposed Zoom-In network consisting on encouraging the model to make predictions for cases ($y = 1$), *e.g.*, images with cancer metastases (see Experiments for details), but using sub-tiles with low attention weights while *inverting* the image labels ($y = 1 \rightarrow 0$). Conveniently, we can generate these (negative) contrastive samples without the need for additional modules or model parameters.

Specifically, we leverage the existing attention functions in (1) and (4). To generate the contrastive feature vectors for image x such that $y = 1$, we first sample (with replacement) tile locations via $1 - a_{\Theta}(V(x, s_1))$ similar to (1). Then, we sample N sub-tiles via $1 - b_{\Theta}(V(x, s_2, c))$ without replacement similar to (4).

The sampled contrastive sub-tiles are passed through the feature network and then processed by the classifier to make predictions $\Psi_{\Theta}(x | y = 1)$ using (9), where the conditioning $y = 1$ is used to emphasize that we use images x of class $y = 1$ as contrastive examples. In general, the number of contrastive examples (per training batch) is equal to the number of samples such that $y = 1$. For these

contrastive sample, we optimize the following objective, $\mathcal{L}_{\text{con}}(\Psi_{\Theta}(x|y=1)) = \sum_n -\log(1 - \Psi_{\Theta}(x_n|y_n=1))$. Note that $\mathcal{L}_{\text{con}}(\Psi_{\Theta}(x|y=1))$ encourages contrastive samples for images x with label $y=1$ to be predicted as $y=0$. In multi-class scenarios, this contrastive learning approach can be readily extended by letting one of the classes be the reference, or in general, by using a complete, cross-entropy-based contrastive loss, in which contrastive samples are generated for both classes, *i.e.*, $y = \{0, 1\}$, instead of just one class (half the cross-entropy loss) as in our case.

3. Related Work

Below we discuss existing research work on classification of very large images with tiny objects, the most relevant body of work being on attention-based models, and general efforts toward computational efficiency for image classification models.

Tiny object classification Recent works have studied CNNs under different noise scenarios, either by performing experiments where label noise is introduced artificially [1, 52], or by directly working with noisy labels and annotations [14, 32]. While it has been shown that large amounts of label noise hinders the generalization ability of CNNs [1, 52], it has been further demonstrated that CNNs can mitigate this label-corrupting noise by increasing the size of the data used for training [32], tuning the parameters of the optimization procedure [19], or re-weighting input training samples [14]. However, all of these works focus on label corruption but do not consider the case of noiseless labels or label assignments with low-noise, in which alternatively, the region of interest (ROI) associated with the label is small or tiny, relative to the size of the image. Purposely, [37] analyzed the capacity of CNNs in precisely this context, *i.e.*, that of tiny object image classification tasks. Their results indicate that by using a training dataset limited in size, CNNs are unable to generalize well as the ROI-to-image ratio of the input decreases. Typically, the object associated with the label occupies a dominant portion of the image. However, in some real-world applications, such as medical imaging, remote sensing or traffic signs recognition, only a very tiny fraction of the image informs their labels, leading to a low ROI-to-image ratios.

Attention This technique has a long history in the neural networks literature [17]. In the modern era of deep learning, it has been used very successfully in various problems [9, 12, 49, 53]. Two main classes of attention mechanisms include: *soft attention*, which estimates a (continuous) weight for every location of the entire input [16], and *hard attention* which selects a fraction of the data, *e.g.*, a ROI in an image, for processing [36], which is a harder problem that resembles object detection, but without ground-truth object boundaries. Note that the attention in [3, 16, 20] is defined as the weights of a bag of features from an image. Our formu-

lation can also be interpreted in the same way, since α is the attention on bag of features of tiles and β^c is the attention on bag of features of sub-tiles.

Computational efficiency There are multiple ways to control the computational cost of deep neural networks. We categorize them into four groups: *i*) compression methods that aim to remove redundancy from already trained models [51]; *ii*) lightweight design strategies used to replace network components with computationally lighter counterparts [18]; *iii*) partial computation methods selectively utilize units of a network, thus creating forward-propagation paths with different computational costs [26]; and *iv*) reinforcement learning and attention mechanisms that can be used to selectively process subsets of the input, based on their importance for the task of interest [7, 20, 28, 40, 47]. The latter being the strategy we consider in the proposed Zoom-In architecture.

In-sample contrastive learning Conventional deep neural networks lack robustness to out-of-distribution data or naturally-occurring corruption such as image noise, blur, compression and label corruption. Contrastive learning [22] has demonstrated great success with learning in noisy scenarios, *e.g.*, corrupted ImageNet [21]. Here, we aim to leverage contrastive learning to mitigate the performance loss caused by images with low ROI-to-image ratio. Hence, the built-in attention mechanism facilitates the in-sample contrastive learning because contrastive samples can be obtained using the same attention mechanisms without the need for additional model components or parameters.

Weakly supervised training on gigapixel images Recent work has demonstrated that deep learning algorithms have the ability to predict patient-level attributes, *e.g.*, cancer staging from whole slide images (WSIs) in digital pathology applications [27]. Because these images are so large and no prior knowledge of which subsets of the image (tiles) are associated with the label, such task is known as weakly supervised learning [4, 31]. Specifically, the model has to estimate which regions within the image are relevant to the label, so predictions can be made using information from these regions alone; not the whole image. Importantly, with current hardware architectures, WSIs are too large to fit in GPU memory, so one commonly used technique is to build a model to select a subset of patches from the image [35, 54]. An alternative approach consists in using the entire WSI but in a compressed, much smaller, representation, at the cost of losing fine-grained details that may be important [45]. Building representations that aggregate features from selected image regions (tiles or patches) are also alternative approaches [4, 8]. We consider the performance of these approaches relative to the proposed Zoom-In network in the experiments.

4. Experiments

We evaluate the proposed approach in terms of accuracy and GPU memory requirements. In the results below, Zoom-In Network refers to the proposed method with a lightweight LeNet backbone, and Zoom-In Network (Res) refers to our method using a ResNet16 backbone. The details of the model architecture are presented in the SM. Moreover, we highlight the ability of the model to attend to a small amount of full-resolution sub-tiles (ROIs) of the image inputs, which results in a significantly reduced peak GPU memory usage and superior test accuracy relative to the competing approaches. We consider methods that can handle large images as inputs, *e.g.*, attention sampling models (ATS) [20], Differentiable Patch Selection (Top-K) [7], BagNet [37], EfficientNet [44] and streaming CNNs [38]. Simultaneously, we also compare our model with methods that apply strategies similar to zoom-in *e.g.*, PatchDrop [47] and RA-CNN [12]. For the peak memory usage, we report inference memory in Mb per sample, *i.e.*, for a batch of size 1. We also report the floating point operations (FLOPs) and run time when inferring a single image. Details of the model architecture and some hyper-parameters not specified in each experiment, as well as an ablation study examining the impact on performance of N , λ , and using contrastive learning are presented in the SM. In-sample contrastive learning is applied after training without it for 10 epochs, and the entropy regularization parameter (for our model and ATS) is set to $\lambda = 1e^{-5}$.

Datasets We focus on datasets that have relatively large image sizes and feature tiny ROI objects that are scattered in large backgrounds, unlike natural images, in which objects are (usually) in the middle of the image due to the fact that photographers tend to center images around the object of interest (target) [46]. Consequently, in the experiments, we do not consider datasets such as ImageNet, iNaturalist and COCO because in these the ROI-to-image ratio is close to 1 and also because image sizes are relatively small relative to some of the other datasets considered and described below.

We present experiments on five datasets: *i*) the colon cancer dataset introduced in [42] aims to detect whether epithelial cells exist in a Hematoxylin and Eosin (H&E) stained images. This dataset contains 100 images of dimensions 500×500 . The images originate both from malignant and normal tissue, and contain approximately 22,000 annotated cells. Following [16, 20], we treat the problem as a binary classification task where the positive images are those containing at least one cell belonging to the epithelial cell class. *ii*) The NeedleCamelyon dataset [37] is built from cropped images from the original Camelyon16 dataset with specified ROI-to-image ratios. Specifically, we generate datasets for ROI-to-image ratios in the range of $[0.1, 1]\%$, and we crop each image with size of $1,024 \times 1,024$ pixels. Positive examples are created by taking 50 random crops

Table 1. Test set results for colon cancer, NeedleCamelyon and fMoW data. Memory denotes average peak memory per sample usage at inference.

Colon Cancer				
Method	Accuracy (%)	FLOPs (B)	Memory (Mb)	Time (ms)
PatchDrop [47]	81.0	75.29	520.44	12.33
RA-CNN [12]	86.4	135.88	4432.46	38.74
CNN [20]	90.8 ± 1.2	1.83	235.68	7.62
ATS [20]	90.7 ± 1.4	0.24	15.83	2.81
Zoom-In Network (ours)	95.0 ± 2.6	0.24	2.55	3.20
NeedleCamelyon				
Method	Accuracy (%)	FLOPs (B)	Memory (Mb)	Time (ms)
BagNet [3]	70.0	222.72	3914.81	12.90
ATS [20]	72.5	1.66	37.97	9.27
Zoom-In Network (ours)	76.0	0.52	11.78	10.20
Zoom-In Network (Res) (ours)	78.1	0.84	14.22	11.42
Traffic Signs Recognition				
Method	Accuracy (%)	FLOPs (B)	Memory (Mb)	Time (ms)
EfficientNet-B0 s0.5 [44]	65.9	4.82	673.39	27.06
EfficientNet-B0 [44]	79.1	19.26	2229.59	34.88
ATS-10 [20]	90.5	1.43	54.51	10.3
TopK-10 [7]	91.7	1.61	125.16	9.8
Zoom-In Network (ours)	91.2	0.79	12.65	12.28
Zoom-In Network (Res) (ours)	92.6	1.18	15.83	13.16
Functional Map of the World				
Method	Accuracy (%)	FLOPs (B)	Memory (Mb)	Time (ms)
EfficientNet-B0 [44]	70.2	8.22	1404.09	22.72
ATS-30 [20]	71.1	2.52	53.42	10.73
TopK-30 [7]	71.6	2.30	73.49	10.12
Zoom-In Network (ours)	72.9	1.85	10.81	11.47
Zoom-In Network (Res) (ours)	74.3	2.24	13.53	12.25

from every annotated cancer metastasis area if the ROI-to-image ratio falls within the range $[0.1, 1]\%$. Negative examples are taken by randomly cropping normal whole-slide images and filtering out image crops that mostly contain background. Further, we ensure the class balance by sampling an equal amount of positive and negative crops. *iii*) Traffic Signs Recognition dataset [25] consists of over 20,000 road scene images with a size of 960×1280 . Here, we use the same subset as [7, 20]. The task is to classify whether the road-scene image contains a speed limit sign (50, 70 or 80 km/h) or not. The subset used in [7, 20] and our experiments includes 747 images for training and 684 images for testing. *iv*) The Functional Map of the World (fMoW) dataset [5] aims to classify the functional purpose of buildings and infrastructures and land-use from high-resolution satellite images. The approximate range of image size in this dataset is 500×500 to $9,000 \times 9,000$ pixels. In our experiment, we extract a class-balanced subset from the original fMoW dataset to further illustrate the proposed method is applicable beyond digital pathology images. Our constructed subset consists of 15,000 training images and 9,571 test image from 10 classes. More details of constructing the subset is provided in SM. *v*) Finally, we utilize the Camelyon16 dataset to further demonstrate the utility of our model on gigapixel images. This dataset contains 400

WSIs, 270 WSIs with pixel-level annotations, and 130 unlabeled WSIs as test set. We split the 270 slides into train and validation sets; for hyperparameter tuning. Typically, only a small portion of a slide contains biological tissue of interest, with background and fat encompassing the remaining areas, *e.g.*, see Figure 1 for a typical WSI (with pixel-level annotations).

Colon cancer We use the same experimental setup as [16, 20], namely, 10-fold-cross-validation, and five repetitions per experiment. The approach most closely related to ours is the one-stage attention sampling model in [20]. We refer to this method as $ATS-N$, where $N = 10$ indicates the number of tiles drawn from the attention weights, and we set $s = s_2$ in all experiments. The value of N was selected to maximize performance. To showcase the advantages of our model compared to traditional CNN approaches, we also include a ResNet [15] with 8 convolutional layers and 32 channels as naive baseline. For our model, we set $s_1 = 0.1$, $s_2 = 0.2$, $N = 10$, and $h_2 = w_2 = 27$.

Results are summarized in Table 1. The proposed two-stage attention sampling model results in approximately 4.3% higher test accuracy than (one-stage) $ATS-10$; presumably due to its ability to better focus on informative regions (sub-tiles) of the image via the hierarchical attention mechanism. Moreover, the baseline (CNN) and $ATS-10$ require at least $\times 90$ and $\times 6$ more memory relative to the Zoom-In network, respectively. Alternatively, Patch-Drop [47] and RA-CNN [12] not only underperform the base CNN but also have higher memory requirements. The memory efficiency of the proposed approach is justified by the way in which the image is processed as a small collection of sub-tiles, thus resulting in a substantially reduced forward pass cost relative to the CNN and $ATS-10$ models. The FLOPs of ATS and Zoom-In Network are comparable because in the colon cancer experiments, the feature extractor $f_{\Theta}(\cdot)$ dominates the FLOP count. Since ATS and Zoom-In Network feed the same number of patches into $f_{\Theta}(\cdot)$, their resulting FLOPs in this step is the same. In fact, this step contributes 0.235 B FLOPs (96%). Further, our model takes slightly more run time due to the implementation inefficiency when extracting tiles and sub-tiles.

NeedleCamelyon Note that the NeedleCamelyon dataset uses larger images ($1,024 \times 1,024$) compared to the ones (up to 512×512) used in [37] to better showcase the abilities of the models considered. Following [37], we split the NeedleCamelyon dataset into training set, validation set, test set with a ratio of 60:20:20. The number of images for each set are 6,000, 2,000 and 2,000, respectively. Positive and negative samples are balanced in each set.

We compare our model with $ATS-30$ and an existing CNN architecture used for a similar NeedleCamelyon dataset in [37]. BagNet [3] is a CNN model extracts features at tile-level, which is efficiently used in NeedleCamelyon




Figure 3. Illustration of the intermediate results for the Traffic Sign dataset using the Zoom-In Network. We show the tile and sub-tile with the highest first-stage and second-stage attention. More intermediate results are included in the SM.

experiments in [37]. We use the BagNet as the CNN baseline in our experiment. For our model, we set $s_1 = 0.25$, $s_2 = 0.5$, $N = 30$, and $h_2 = w_2 = 32$. Due to the much smaller ROI-to-image ratio of NeedleCamelyon relative to the colon cancer dataset, we set a larger s_1 and s_2 to ensure the down-sampling will not wash out the discriminative information. We also increase N since the number of attention weights with large values is larger in this case.

Table 1 shows performance for the proposed model and the baselines. We observe that for larger images, the Zoom-In network has better GPU memory use at inference time because much less sub-tiles at scale s_2 tend to be instantiated. In terms of test accuracy, the proposed model results in approximately 3.5% higher test accuracy than one stage $ATS-30$ and 6.0% higher than the BagNet baseline in this experiment. Moreover, BagNet and one-stage attention sampling require at least $\times 500$ and $\times 7$ more memory, respectively, compared to the proposed approach. Here, we can see our Zoom-In Network consumes drastically less FLOPs than ATS . This is because the Zoom-In Network requires less pixels of the original input images to select a comparable amount of high resolution ROI patches for prediction relative to ATS .

Traffic Sign Recognition We compare the Zoom-In Network with a traditional CNN (EfficientNet-B0 [44]) and the recently published one-stage zoom-in methods (ATS [20], TopK [7]). For EfficientNet-B0, we use both the original resolution images and images downsampled by half (denoted as $s0.5$) as inputs, to show the limitations of traditional CNNs. For ATS and TopK, we use the same hyperparameter settings from [7, 20]. For our Zoom-In Network, we tried two types of the backbones (LeNet and ResNet16) and we set $s_1 = 0.125$, $s_2 = 0.3$, $N = 10$, and $h_2 = w_2 = 100$. The details of the network architecture are shown in the SM.

In Table 1, we can see that our Zoom-In model achieves the highest accuracy and lowest GPU memory consumption among all tested methods. Also, our model requires less FLOPs by using less number of input pixels than ATS and Top-K. The merits of our model originate from accurate target object localization achieved by the attention mechanism and contrastive learning, as well as efficient GPU memory usage. Examples of attention maps for Traffic Sign dataset similar to those in Figure 3 are shown in the SM.

Table 2. Results on Camelyon16 data. Memory denotes average peak memory per sample at inference.

Method	Pixel-level Annotation	Accuracy (%)	Memory (Mb)
Streaming CNN [38]	No	70.6	3,256.29
CLAM [31]	No	78.0	206.88
MIL [4]	No	79.9	140.68
MRMIL [29]	No	81.1	568.00
Zoom-In Network	No	81.3	71.76
Zoom-In Network (Res)	No	82.6	71.76
Zoom-In Network	Yes	88.2	71.76
Zoom-In Network(Res)	Yes	90.8	71.76
Winning model [2]	Yes	92.2	395.77

Functional Map of the World Functional Map of the World (fMoW) consists of a large amount of high-resolution RGB images of various sizes ranging from 500×500 to $9,000 \times 9,000$. Following [36], we choose EfficientNet-B0 [44] as the baseline model. EfficientNet-B0 effectively scales with large images and has been proved to serve as a good baseline model on the fMoW dataset in [36]. We also implement the ATS and TopK with our best efforts for fMoW dataset. We set $s_1 = 0.25$, $s_2 = 0.5$, $N = 30$, and $h_2 = w_2 = 50$ for our model.

Results in Table 1 show that Zoom-In Network surpasses EfficientNet-B0 in accuracy and memory consumption at inference time. Examples of attention maps for fMoW similar to those in Figure 3 are shown in the SM.

Camelyon16 This is a gigapixel dataset consisting of whole-slide images (WSIs) with sizes ranging from $45,056 \times 35,840$ to $217,088 \times 111,104$. The objective here is to predict whether a WSI contains cancer metastases. As we described in the Related Work Section, existing works have attempted to train CNNs on very large images with only image-level labels, *i.e.*, via weakly supervised training on gigapixel images. We consider the streaming CNN [38], CLAM [31], MIL [4] and MRMIL [29], which we previously briefly described. Details of these baselines are presented in the SM. Further, we also evaluate the model for the scenario in which pixel-level annotations (ROIs) are available. Here, we compare our results to the winning model of the Camelyon16 challenge [2]. For our model, we set $s_1 = 0.03125$, $s_2 = 0.125$, $N = 100$, and $h_2 = w_2 = 50$.

In Table 2, we see that our Zoom-In network achieves the highest test accuracy when pixel-level annotations are not available. Even when there are pixel-level annotations, our model yields a test accuracy close to the winning model of the Camelyon16 challenge, without the need for comprehensive tuning and handcrafted features (*i.e.*, the minimum surrounding convex region, the average prediction values, and the longest axis of the lesion area). For the memory comparison, all methods require substantially more memory than the proposed approach, notable $\times 8$ more than MRMIL which has comparable performance. Moreover, we also consider the situation in which the pixel-level anno-

tations are available. The extension to leverage pixel-level annotations is described in the SM. The performance of the proposed model is relatively close to the winning model [2], which indicates that the proposed approach is flexible and also accurate when we have the manually annotated ROIs. Additional details including attention maps similar to that in Figure 3 are provided in the SM.

We also analyzed the correlation of the attention weights generated by the Zoom-In network with the ground-truth metastases-to-tile ratio obtained from pixel annotations. Specifically, the proportion of each sub-tile of size $h_2 \times w_2$ that is covered by cancer metastases. The Spearman correlation coefficient for all the tiles with pixel-level annotations is $\rho = 0.3570$, indicating a good agreement. Note that these attention weights are obtained from the model trained without pixel-level annotation information. In the SM, we visually present these correlations as scatter plots (attention weights *vs.* metastases-to-tile ratios).

From all the results above, we see Zoom-In Network reduces the GPU memory usage significantly. The reasons why we pursue a low GPU memory consumption are: *i)* GPUs with high memory are expensive and not widely available for applications in practice. A model using less GPU at inference time can be deployed with less expense; *ii)* the highly memory-efficient model allows training and inferring images larger than gigapixels possible; *iii)* With the growing use of neural network on mobile/edge devices, it is key to develop memory-light GPU models to allow locally running deep learning service on mobile/edge devices.

5. Discussion

We presented the Zoom-In network that can efficiently classify very large images with tiny objects. We improved over the existing CNN-based baselines both in terms of accuracy and peak GPU memory use, by leveraging a two-stage hierarchical sampling strategy and a contrastive learning objective. We also considered the scenario in which pixel-level annotations (segmentation maps) are available during training. In the experiments, we demonstrated the advantage of the proposed model on five challenging classification tasks. We note that the images in the Camelyon16 dataset are all gigapixel in size and that we are likely the first ones training an end-to-end deep learning model for them. Our model achieved the best accuracy when pixel-level annotations are not available, while also using a small amount of GPU memory, which allows for training and inference on full-resolution gigapixel images using a single GPU. One limitation of the proposed model is the need to specify the number of sub-tile samples N , which can be potentially estimated from data.

Acknowledgements This work was supported by NIH (R44-HL140794), DARPA (FA8650-18-2-7832-P00009-12) and ONR (N00014-18-1-2871-P00002-3).

References

- [1] Devansh Arpit, Stanisław Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International Conference on Machine Learning*, pages 233–242. PMLR, 2017. 5
- [2] Babak Ehteshami Bejnordi, Mitko Veta, Paul Johannes Van Diest, Bram Van Ginneken, Nico Karssemeijer, Geert Litjens, Jeroen AWM Van Der Laak, Meyke Hermesen, Quirine F Manson, Maschenka Balkenhol, et al. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *JAMA*, 318(22):2199–2210, 2017. 8
- [3] Wieland Brendel and Matthias Bethge. Approximating cnns with bag-of-local-features models works surprisingly well on imagenet. *arXiv preprint arXiv:1904.00760*, 2019. 5, 6, 7
- [4] Gabriele Campanella, Matthew G Hanna, Luke Geneslaw, Allen Miraflor, Vitor Werneck Krauss Silva, Klaus J Busam, Edi Brogi, Victor E Reuter, David S Klimstra, and Thomas J Fuchs. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature medicine*, 25(8):1301–1309, 2019. 5, 8
- [5] Gordon Christie, Neil Fendley, James Wilson, and Ryan Mukherjee. Functional map of the world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6172–6180, 2018. 2, 6
- [6] Robert Graham Clark. Sampling of subpopulations in two-stage surveys. *Statistics in Medicine*, 28(29):3697–3717, 2009. 3
- [7] Jean-Baptiste Cordonnier, Aravindh Mahendran, Alexey Dosovitskiy, Dirk Weissenborn, Jakob Uszkoreit, and Thomas Unterthiner. Differentiable patch selection for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2351–2360, 2021. 5, 6, 7
- [8] Pierre Courtiol, Eric W Tramel, Marc Sanselme, and Gilles Wainrib. Classification and disease localization in histopathology using only global labels: A weakly-supervised approach. *arXiv preprint arXiv:1802.02212*, 2018. 5
- [9] Misha Denil, Loris Bazzani, Hugo Larochelle, and Nando de Freitas. Learning where to attend with deep architectures for image tracking. *Neural computation*, 24(8):2151–2184, 2012. 5
- [10] Xuanyi Dong, Junshi Huang, Yi Yang, and Shuicheng Yan. More is less: A more complicated network with less inference complexity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5840–5848, 2017. 2
- [11] Pavlos S Efraimidis and Paul G Spirakis. Weighted random sampling with a reservoir. *Information Processing Letters*, 97(5):181–185, 2006. 4
- [12] Jianlong Fu, Heliang Zheng, and Tao Mei. Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4438–4446, 2017. 2, 5, 6, 7
- [13] Lindsay P Galway, Nathaniel Bell, Sahar AE Al Shatari, Amy Hagopian, Gilbert Burnham, Abraham Flaxman, William M Weiss, Julie Rajaratnam, and Tim K Takaro. A two-stage cluster sampling method using gridded population data, a gis, and google earth tm imagery in a population-based mortality survey in iraq. *International Journal of Health Geographics*, 11(1):1–9, 2012. 3
- [14] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *arXiv preprint arXiv:1804.06872*, 2018. 5
- [15] 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. 7, 6
- [16] Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International Conference on Machine Learning*, pages 2127–2136. PMLR, 2018. 5, 6, 7
- [17] Laurent Itti, Christof Koch, and Ernst Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11):1254–1259, 1998. 5
- [18] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. *arXiv preprint arXiv:1405.3866*, 2014. 5
- [19] Stanisław Jastrzebski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three factors influencing minima in SGD. *arXiv preprint arXiv:1711.04623*, 2017. 5
- [20] Angelos Katharopoulos and François Fleuret. Processing megapixel images with deep attention-sampling models. In *International Conference on Machine Learning*, pages 3282–3291. PMLR, 2019. 2, 3, 4, 5, 6, 7
- [21] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*, 2020. 5
- [22] Minsong Ki, Youngjung Uh, Wonyoung Lee, and Hyeran Byun. In-sample contrastive learning and consistent attention for weakly supervised object localization. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 4, 5
- [23] Wouter Kool, Herke Van Hoof, and Max Welling. Stochastic beams and where to find them: The gumbel-top-k trick for sampling sequences without replacement. In *International Conference on Machine Learning*, pages 3499–3508. PMLR, 2019. 4
- [24] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25:1097–1105, 2012. 1
- [25] Fredrik Larsson and Michael Felsberg. Using fourier descriptors and spatial models for traffic sign recognition. In

- Scandinavian conference on image analysis*, pages 238–249. Springer, 2011. 6
- [26] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Ultra-deep neural networks without residuals. In *Int. Conf. on Learning Representations, arXiv, Toulon, France*, page 1605, 2017. 5
- [27] Byungjae Lee and Kyunghyun Paeng. A robust and effective approach towards accurate metastasis detection and pn-stage classification in breast cancer. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 841–850. Springer, 2018. 5
- [28] Hila Levi and Shimon Ullman. Efficient coarse-to-fine non-local module for the detection of small objects. *arXiv preprint arXiv:1811.12152*, 2018. 5
- [29] Jiayun Li, Wenyuan Li, Anthony Sisk, Huihui Ye, W Dean Wallace, William Speier, and Corey W Arnold. A multi-resolution model for histopathology image classification and localization with multiple instance learning. *Computers in Biology and Medicine*, 131:104253, 2021. 8
- [30] Geert Litjens, Peter Bandi, Babak Ehteshami Bejnordi, Oscar Geessink, Maschenka Balkenhol, Peter Bult, Altuna Halilovic, Meyke Hermsen, Rob van de Loo, Rob Vogels, et al. 1399 h&e-stained sentinel lymph node sections of breast cancer patients: the camelyon dataset. *GigaScience*, 7(6):giy065, 2018. 1, 2
- [31] Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Richard J Chen, Matteo Barbieri, and Faisal Mahmood. Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature Biomedical Engineering*, pages 1–16, 2021. 5, 8
- [32] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens Van Der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European Conference on Computer Vision*, pages 181–196, 2018. 5
- [33] Francesco Marra, Diego Gragnaniello, Luisa Verdoliva, and Giovanni Poggi. A full-image full-resolution end-to-end-trainable cnn framework for image forgery detection. *IEEE Access*, 8:133488–133502, 2020. 2
- [34] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pages 1928–1937. PMLR, 2016. 6
- [35] Nikhil Naik, Ali Madani, Andre Esteva, Nitish Shirish Keskar, Michael F Press, Daniel Ruderman, David B Agus, and Richard Socher. Deep learning-enabled breast cancer hormonal receptor status determination from base-level h&e stains. *Nature Communications*, 11(1):1–8, 2020. 5
- [36] Athanasios Papadopoulos, Paweł Korus, and Nasir Memon. Hard-attention for scalable image classification. *arXiv preprint arXiv:2102.10212*, 2021. 2, 5, 8, 4
- [37] Nick Pawlowski, Suvrat Bhooshan, Nicolas Ballas, Francesco Ciompi, Ben Glocker, and Michal Drozdal. Needles in haystacks: On classifying tiny objects in large images. *arXiv preprint arXiv:1908.06037*, 2019. 1, 5, 6, 7, 4
- [38] Hans Pinckaers, Bram van Ginneken, and Geert Litjens. Streaming convolutional neural networks for end-to-end learning with multi-megapixel images. *arXiv preprint arXiv:1911.04432*, 2019. 2, 6, 8
- [39] Suo Qiu. Global weighted average pooling bridges pixel-level localization and image-level classification. *arXiv preprint arXiv:1809.08264*, 2018. 7
- [40] Jason Ramapuram, Maurits Diephuis, Frantzeska Lavda, Russ Webb, and Alexandros Kalousis. Variational saccading: Efficient inference for large resolution images. *arXiv preprint arXiv:1812.03170*, 2018. 5
- [41] Minsoo Rhu, Natalia Gimelshein, Jason Clemons, Arslan Zulfiqar, and Stephen W Keckler. vdn: Virtualized deep neural networks for scalable, memory-efficient neural network design. In *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 1–13. IEEE, 2016. 4
- [42] Korsuk Sirinukunwattana, Shan E Ahmed Raza, Yee-Wah Tsang, David RJ Snead, Ian A Cree, and Nasir M Rajpoot. Locality sensitive deep learning for detection and classification of nuclei in routine colon cancer histology images. *IEEE Transactions on Medical Imaging*, 35(5):1196–1206, 2016. 6
- [43] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Mask-guided contrastive attention model for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1179–1188, 2018. 7
- [44] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019. 1, 6, 7, 8
- [45] David Tellez, Geert Litjens, Jeroen van der Laak, and Francesco Ciompi. Neural image compression for gigapixel histopathology image analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 5
- [46] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. Fixing the train-test resolution discrepancy. *arXiv preprint arXiv:1906.06423*, 2019. 6
- [47] Burak Uzsent and Stefano Ermon. Learning when and where to zoom with deep reinforcement learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12345–12354, 2020. 5, 6, 7
- [48] Dayong Wang, Aditya Khosla, Rishab Gargeya, Humayun Irshad, and Andrew H Beck. Deep learning for identifying metastatic breast cancer. *arXiv preprint arXiv:1606.05718*, 2016. 7
- [49] Zhe Wang, Yanxin Yin, Jianping Shi, Wei Fang, Hongsheng Li, and Xiaogang Wang. Zoom-in-net: Deep mining lesions for diabetic retinopathy detection. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 267–275. Springer, 2017. 5
- [50] Wei Xia, Caihong Ma, Jianbo Liu, Shibin Liu, Fu Chen, Zhi Yang, and Jianbo Duan. High-resolution remote sensing imagery classification of imbalanced data using multistage sampling method and deep neural networks. *Remote Sensing*, 11(21):2523, 2019. 3
- [51] Ruichi Yu, Ang Li, Chun-Fu Chen, Jui-Hsin Lai, Vlad I Morariu, Xintong Han, Mingfei Gao, Ching-Yung Lin, and

- Larry S Davis. NISP: Pruning networks using neuron importance score propagation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9194–9203, 2018. 5
- [52] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016. 5
- [53] Heliang Zheng, Jianlong Fu, Zheng-Jun Zha, and Jiebo Luo. Looking for the devil in the details: Learning trilinear attention sampling network for fine-grained image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5012–5021, 2019. 5
- [54] Xinliang Zhu, Jiawen Yao, and Junzhou Huang. Deep convolutional neural network for survival analysis with pathological images. In *2016 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 544–547. IEEE, 2016. 5