

RADIOLOGIST-LEVEL COVID-19 DETECTION USING CT SCANS WITH DETAIL-ORIENTED CAPSULE NETWORKS

APRIL 17, 2020

Aryan Mobiny^{1,*†}, Pietro A. Cicalese^{1,2,†}, Samira Zare¹, Pengyu Yuan¹, Mohammad S. Abavisan¹, Carol C. Wu³, Jitesh Ahuja³, Patricia M. de Groot³, and Hien V. Nguyen¹

¹Department of Electrical and Computer Engineering, University of Houston, Houston, TX, 77004 USA

²Department of Biomedical Engineering, University of Houston, Houston, TX, 77004 USA

³Department of Thoracic Imaging, University of Texas MD Anderson Cancer Center, Houston, TX, 77030 USA

[†]These authors contributed equally to the completion of this work.

*corresponding author: amobiny@uh.edu

ABSTRACT

Radiographic images offer an alternative method for the rapid screening and monitoring of Coronavirus Disease 2019 (COVID-19) patients. This approach is limited by the shortage of radiology experts who can provide a timely interpretation of these images. Motivated by this challenge, our paper proposes a novel learning architecture, called Detail-Oriented Capsule Networks (DECAPS), for the automatic diagnosis of COVID-19 from Computed Tomography (CT) scans. Our network combines the strength of Capsule Networks with several architecture improvements meant to boost classification accuracies. First, DECAPS uses an Inverted Dynamic Routing mechanism which increases model stability by preventing the passage of information from non-descriptive regions. Second, DECAPS employs a Peekaboo training procedure which uses a two-stage patch crop and drop strategy to encourage the network to generate activation maps for every target concept. The network then uses the activation maps to focus on regions of interest and combines both coarse and fine-grained representations of the data. Finally, we use a data augmentation method based on conditional generative adversarial networks to deal with the issue of data scarcity. Our model achieves 84.3% precision, 91.5% recall, and 96.1% area under the ROC curve, significantly outperforming state-of-the-art methods. We compare the performance of the DECAPS model with three experienced, well-trained thoracic radiologists and show that the architecture significantly outperforms them. While further studies on larger datasets are required to confirm this finding, our results imply that architectures like DECAPS can be used to assist radiologists in the CT scan mediated diagnosis of COVID-19.

Keywords AI Diagnosis, Capsule Networks, Coronavirus, COVID-19, Computed Tomography, Visualization.

1 INTRODUCTION

The Coronavirus Disease 2019 (COVID-19) pandemic has taken more than 100,000 human lives, has cost the world economy several trillion dollars, and has fundamentally changed every aspect of our lives. Early, accurate, and scalable technologies for the screening of patients and monitoring of their treatment progress are crucial for controlling this crisis. At the time of writing this paper, real-time Reverse Transcriptase Polymerase Chain Reaction (rRT-PCR) is the most widely used approach for patient screening [1]. Patient specimens are usually collected from their nose or throat, and are then sent to laboratories to generate a diagnosis. The process is time-consuming, complicated, and limited by the supply of test kits, swabs, and test reagents, which has led to logistic challenges due to the large influx of patients.

Computed tomography (CT) of the thorax is a tool frequently used to assess patients with fever and/or respiratory symptoms, and it has been used during this pandemic to triage suspected COVID-19 patients for further testing or hospital admission and to detect treatment response or complications [2, 3]. Recent studies show that CT scans of COVID-19 patients have typical features that can be recognized by doctors and artificial intelligence models [4, 2, 5]. While the role of radiographic images is still under discussion as the crisis unfolds [6], they have several

important advantages including fast turnaround time (e.g., 15 to 30 minutes for a CT scan) and wide availability at most hospitals. These advantages potentially make radiographic imaging techniques a great complementary tool to rRT-PCR testing with which to fight the COVID-19 pandemic [7, 8].

An important challenge in using radiographic imaging is the shortage of well-trained radiologists who can provide a timely and accurate interpretation of patients' images [9]. A new disease like COVID-19 with variable imaging manifestations requires radiologists to spend significant amounts of time to constantly update their knowledge and learn new interpretation skills. This becomes even more challenging in remote locations with insufficient access to training resources and poorly-equipped medical facilities. The inability to perform image interpretation for the diagnosis of the disease and assessment of disease severity make these remote regions of the country highly vulnerable when the disease strikes. In addition, the pandemic creates a large influx of patients which sharply increases the radiologists' workload. Addressing the shortage of radiology experts is critical to the diagnostic process and is needed to maximize the potential of radiographic imaging in the diagnosis of COVID-19.

Motivated by the urgent need, recent work has developed a number of artificial intelligence models for automatic diagnosis or

assessment of COVID-19 from CT scans and chest radiographic data [10, 11, 12, 13]. Unfortunately, training these models requires a large number of annotated examples. This is a difficult requirement given that data are not abundantly available nor properly annotated during the initial period of the pandemic, partly because physicians are too busy to provide clinical feedback. In addition, most prior publications do not make their data publicly available due to privacy and proprietary issues [14, 2, 11]. Therefore, being able to build robust and accurate artificial intelligence models using a small number of imperfect samples is highly desirable.

Existing work is limited to evaluating artificial intelligence models using classification accuracy on an independent test set. While the results are encouraging, it remains unclear how these models fare against well-trained human radiologists and how they derive their conclusions. Past work with deep learning has demonstrated that deep networks are competitive against human experts on a number of classification tasks for chest radiographs [15, 16], CT scans [17], and magnetic-resonance images [18]. However, the nature of these studies is significantly different from the one in this paper. In particular, while these studies use large-scale and carefully curated datasets, our study focuses on building artificial intelligence models from a modest and possibly noisy set of training examples. In addition, the pathological conditions in prior studies are well-established while COVID-19 is still a new disease whose characteristics have not been fully understood. Comparing the performance of artificial intelligence algorithms against human experts at the developing stage of a new disease is unique. It can potentially shed light on how well computer and human experts stand up to a new diagnostic challenge. Our results are expected to provide a more complete view of artificial intelligence algorithms, including their strengths and weaknesses in responding to the COVID-19 pandemic. Our paper makes the following contributions:

- We propose a novel detail-oriented capsule network architecture capable of identifying fine-grained and discriminative image features to classify COVID-19 patients.
- Our model employs a data augmentation method based on conditional generative adversarial nets that is meant to deal with the issue of data scarcity.
- We provide a rigorous evaluation of our models with visualization and clinical input from clinicians. Our extensive experiments on the COVID-19 CT dataset indicate that the proposed models are capable of learning meaningful and discriminative features.
- Our paper compares the performances of artificial intelligence models to three well-trained and experienced radiologists. To the best of our knowledge, this is the first attempt to benchmark deep learning models against human experts for COVID-19 diagnosis.

The rest of this paper is organized as follows: works related to automatic diagnosis of COVID-19, capsule networks, and data augmentation with generative models are presented in Section 2. Section 3 explains the proposed network architecture and the methods applied to address data scarcity. Section 4 describes the dataset used in this study. Experimental results are presented in Section 5 and are discussed in Section 6. Section 7 concludes the

paper with future research directions. All relevant source code and radiologist diagnostic data will be made publicly available.

2 RELATED WORK

Automatic Diagnosis of COVID-19: COVID-Net was recently introduced in [10] for the automatic interpretation of chest radiographs of COVID patients. The network consists of convolutional layers with skip connections similar to residual networks [19]. Authors showed promising accuracy in classifying X-ray images aggregated from various online sources. [20]. Chaganti *et al.* [12] proposed a system that finds suspicious lung regions in a CT image using a deep reinforcement learning algorithm and then measured the abnormality and severity caused by the virus. Li *et al.* [11] modified a residual network with 50 layers to detect COVID-19 using CT scans. The model was trained on a dataset comprised of 4,356 CT scans collected from multiple hospitals in China. The dataset, however, has not been made available to the public. [13] proposed an architecture based on ResNet18 to extract features from CT scans, segment out pathological lung regions which then serve as the input to a classifier for predict COVID-19 disease. The network achieves an overall accuracy of 86.7%. While these studies demonstrated the great potential of artificial intelligence algorithms in automating COVID-19 diagnosis, most of them do not address the problem of data scarcity in model training, which is a practical constraint at the beginning of a pandemic. More importantly, their evaluations are limited to standard classification metrics on independent test sets. None of the work has compared the performances of artificial intelligence models to radiology experts. Such a comparison coupled with the qualitative visualization of our results will improve our understanding of the strengths and weaknesses of our algorithms while ensuring that high accuracies are not the result of over-fitting.

Capsule Networks: Capsule Networks (CapsNets) [21, 22] have demonstrated to be a promising alternative to convolutional neural networks (CNNs). They are capable of encoding part-whole relationships and the relative pose of objects with respect to the surrounding scene [21]. CapsNets employ an iterative routing-by-agreement mechanism to determine where to send information which allows capsules to learn viewpoint invariant representations of the data. CapsNets achieved state-of-the-art performance in many tasks including hand-written digit recognition [21], human action detection in videos [23], and cancer screening with volumetric radiology data [24]. However, CapsNets with dynamic or EM routing experience unstable training when the number of layers increases [25, 26]. It is also unclear how CapsNets select the informative regions within the image to make their decision. Our paper introduces a stable detail-oriented CapsNet architecture that intelligently filters out irrelevant information, effectively capturing both coarse and fine-grained representations within the data.

Generative Models for Data Augmentation: Acquiring a large amount of training samples that effectively characterize population diversity is necessary for the success of deep neural networks in medical applications. However, COVID-19 imaging data availability, especially during a crisis like the COVID-19 pandemic, is quite limited; high annotation costs, overwhelmed physicians, and time restrictions all contribute to data scarcity. Therefore, generating descriptive and realistic synthetic training

samples serves as an enticing solution. Jin *et al.* [27] employed a 3D generative adversarial network to artificially generate fake lung CT scans with lung nodules of various dimensions at multiple locations and achieved promising results. Han *et al.* [28] proposed a 3D multi-conditional GAN (3D MCGAN) which uses context and nodule discriminators with distinct loss functions to generate realistic and diverse nodules in CT images. Ghorbani *et al.* [29] proposed DermGAN (implementing the pix2pix architecture), which can translate skin conditions to a synthetic sample with a new size, location and skin color. We propose the use of the pix2pix architecture as a means of augmenting our relatively scarce COVID-19 2D CT dataset due to the quality, high sample stochasticity, and feature fidelity it achieves.

3 METHODOLOGY

The accurate assessment of a radiology image depends on the physician’s ability to focus on relevant information while simultaneously considering the context of neighboring regions and the entire scan. It would therefore be preferable to simulate this behavior by using a classification architecture that is capable of deriving conclusions about regions of interest with respect to the greater context of the image. To accomplish this, we propose the Detail Oriented Capsule Network (DECAPS) architecture, which uses two underlying concepts to replicate how a radiologist reviews each scan. First, we use a Capsule Network (CapsNet) classifier, which allows the architecture to intelligently combine information within and around a region of interest to generate a holistic interpretation of each scan. Second, we introduce the “peekaboo” training procedure which randomly retrieves or removes areas of interest in the image. When we pass the retrieved ROI and the entire scan through the classifier, the *local* classification power of the architecture is effectively increased by simulating the ability of the radiologist to consider the ROI with respect to the entire scan. Alternatively, we pass the whole scan with the ROI removed; this forces the classifier to recognize multiple ROIs when possible, thus enhancing its *global* classification power.

Radiology data is often comparatively scarce in the context of classification tasks; this is especially true for images derived from novel disease cases like COVID-19. This presents a unique challenge for researchers in how to best utilize their available data to generate robust and reliable models that are able to generalize to the patient population. While popular online augmentation techniques such as image flipping, rotation, and cropping help improve generalization, they fail to minimize the impact of otherwise irrelevant information that is coincidentally correlated to the classification labels. Unique image characteristics such as organ shapes, scan irregularities, and other morphologically irrelevant information can easily distract a classifier trained on a small sample of data, thus harming its ability to generalize. Considering the urgent need for the dependable and high-throughput diagnosis of COVID-19, we believe that leveraging intelligent data augmentation techniques that learn to identify and combine class specific features is fundamental to building an inexpensive but robust COVID-19 classifier. To accomplish this, we propose the use of conditional GANs as a means to model the distribution of the CT image data for each class; this allows us to then generate unique samples from the distribution that can then be

passed to our desired classification architecture to improve its generalization ability.

Notations: Throughout the paper, r , \mathbf{r} , \mathbf{R} , \mathbf{R} represent a scalar, a vector, a 2D matrix, and a tensor (i.e. a higher dimensional matrix; usually a 3D matrix of capsule activations), respectively. Multiplying a transformation matrix and a tensor of poses is equivalent to applying the transformation to each pose.

3.1 Background on Capsule Networks

A CapsNet is composed of a sequence of capsule layers, each of which contains multiple discriminative elements called capsules. We can define each capsule as being a group of neurons that effectively encode a pose vector or matrix [21, 22]. Suppose that we have a set of capsules in layer L which we denote as Ω_L . We then say that each capsule $i \in \Omega_L$ outputs a pose vector \mathbf{p}_i^L ; each element in the matrix characterizes the instantiation parameters (such as orientation, size, etc.). The activation probability of a capsule a_i^L indicates the presence of an entity and is implicitly encoded in the capsule as the Frobenius norm of the pose vector. We then generate a vote vector \mathbf{v}_{ij}^L by processing the information from the i -th capsule in Ω_L to the j -th capsule in Ω_{L+1} by using a linear transformation $\mathbf{v}_{ij}^L = \mathbf{W}_{ij}^L \mathbf{p}_i^L$. The pose of a given capsule $j \in \Omega_{L+1}$ is generated using a convex combination of all the votes from child capsules: $\mathbf{p}_j^{(L+1)} = \sum_i r_{ij} \mathbf{v}_{ij}^L$, where r_{ij} are routing coefficients and $\sum_i r_{ij} = 1$. We generate these coefficients through a dynamic routing algorithm which increases the routing coefficient r_{ij} when the corresponding vote vector \mathbf{v}_{ij}^L and \mathbf{p}_j^{L+1} are similar to each other [21]. This process weights the output of each child capsule to be passed to the corresponding parent capsules. The network therefore constructs a transformation matrix for all capsule pairs to encode the part-whole relationships (i.e. parsing a semantic tree between capsules) while retaining the geometric information of the input data. This property has prompted several research groups to develop new capsule designs and routing algorithms [24, 30, 25, 31].

3.2 Detail Oriented Capsule Networks (DECAPS)

As mentioned in the previous section, the pose of parent capsules within the vanilla CapsNet architecture are derived from the votes of all corresponding children capsules. This can have a negative effect of the performance of the classifier; while some child capsules pass useful information from regions of interest, others pass noise derived from non-descriptive areas. To address this limitation, we implemented a unique CapsNet architecture, loss function, and inverted routing mechanism which increases the weight of votes derived from ROIs. This effectively improves the quality of the input information being passed to each parent, thus strengthening the networks ability to build part-whole relationships. Note that given the relatively small sample size of our dataset, we chose to utilize pose matrices to decrease the number of training parameters to prevent the early over-fitting that is associated with a large number of parameters [22]. We also took inspiration from the Transformers architecture described by Vaswani *et al.* and grouped capsules into grids called *Capsule Heads*, as shown in Fig. 1 (a). Each capsule head is designed to route information to parent capsules independently of the other capsule heads; we encourage this by forcing each head to share a transformation matrix \mathbf{W}_{ij}^L between all capsules for each output

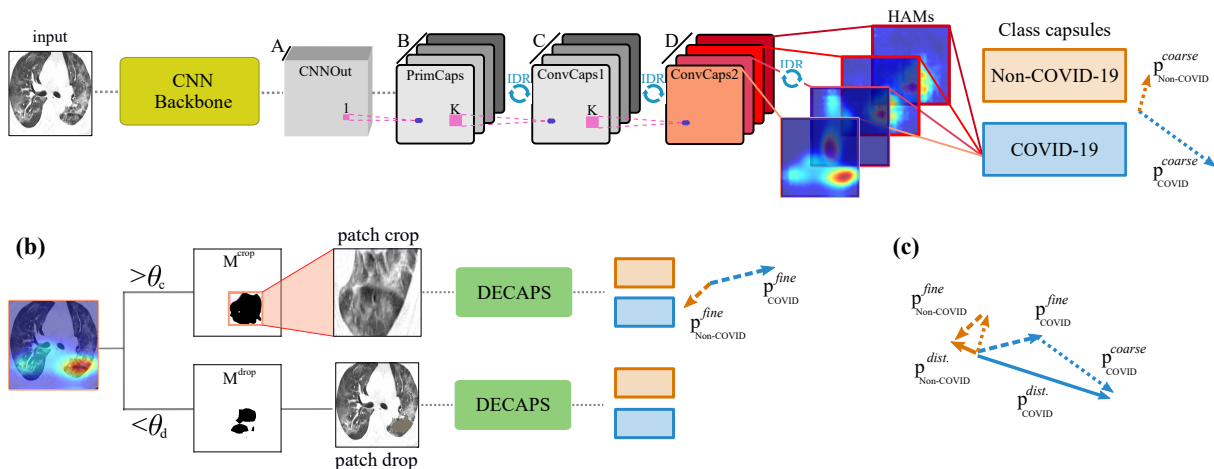


Figure 1: **(a)**: Illustration of the DECAPS architecture. Sample head activation maps (HAMS) are presented for the COVID-19 class. **(b)**: Peekaboo training. HAMS are randomly selected to perform patch crop/drop while training. At test time, the average HAM across all heads is used to perform patch crop (with no patch drop required at test time) to obtain the fine-grained prediction. **(c)**: The distillation process to fine-tune the coarse-grained prediction (p^{coarse}) using the fine-grained prediction (p^{fine}) to obtain the final distilled predictions (p^{dist}).

class. This differs from the vanilla CapsNet architecture which uses one transformation matrix per capsule per class, and effectively reduces the number of trainable parameters by an order of head size (or the number of capsules within a head). This allows us to both train our network with large images and increase the number of capsule layers used (thus increasing abstraction and the complexity of the features captured by the classifier).

We let $\mathbf{P}_i^L \in \mathbb{R}^{h_L \times w_L \times d_L}$ represent the pose matrix of the capsules of the i^{th} head where h_L and w_L represent the height and width of the head respectively, while d_L denotes the capsule dimension (i.e. the number of hidden units grouped together to yield the capsules in layer L). We change i to represent the index of the head in our architecture as opposed to its previous definition in our vanilla CapsNet description (the child capsule index). We take the Frobenius norm of each capsule pose matrix to represent the existence probability of the desired semantic category; these values are squashed through a non-linear squash function to ensure that they fall between zero (not present) and one (present) [21]. We then define the votes from the capsules of the i^{th} head to the j^{th} parent capsule as $\mathbf{V}_{ij}^L = \mathbf{W}_{ij}^L \mathbf{P}_i^L$. For each coordinate within a capsule head, we take the relative coordinates of each capsule (row and column) and add them to the final two entries of the vote vector, allowing us to preserve the capsules location [22]. At this point, the generated votes are processed by the inverted routing mechanism to effectively parse each semantic tree present in the data.

Inverted Dynamic Routing: Dynamic routing [21] is a *bottom-up* approach which forces higher-level capsules to compete with each other to collect lower-level capsule votes. We propose an inverted dynamic routing (IDR) technique which implements a *top-down* approach, effectively forcing lower-level capsules to compete for the attention of higher-level capsules (see Fig. 3.2). During each iteration of the routing procedure, we use a softmax function to force the routing coefficients between all capsules of a single head and a single parent capsule to sum to

one (see Algorithm 1). The pose matrix of the j^{th} parent capsule, \mathbf{P}_j^{L+1} , is then set to the squashed weighted-sum over all votes from the earlier layer (line 6 in Algorithm 1). Given the vote map computed as $\mathbf{V}_{ij}^L = \mathbf{W}_{ij}^L \mathbf{P}_i^L \in \mathbb{R}^{h_L \times w_L \times d_{L+1}}$, the proposed algorithm generates a routing map $\mathbf{R}_{ij}^L \in \mathbb{R}^{h_L \times w_L}$ from each capsule head to each output class. The voting map describes the children capsules' votes for the parent capsule's pose. The routing map depicts the weights of the children capsules according to their agreements with parent capsules, with winners having the highest r_{ij}^L . We combine these maps to generate head activation maps (or HAMS) following

$$A_{ij}^L = \left(\sum_d (\tilde{A}_{ij}^L)^2 \right)^{1/2}, \quad \text{where } \tilde{A}_{ij}^L = \mathbf{R}_{ij}^L \odot \mathbf{V}_{ij}^L \quad (1)$$

in which A_{ij}^L is the head activation map from the i^{th} head in layer L to the j^{th} head in layer $L+1$, and \sum_d is the sum over d_{L+1} channels along the third dimension of \mathbf{V}_{ij}^L . In the output of the last convolutional capsule layer, A_{ij}^L highlights the informative regions within an input image corresponding to the j^{th} class, captured by the i^{th} head. IDR returns as many activa-




Figure 2: Dynamic (left) vs. Inverted dynamic routing (right). Capsule heads are shown in different colors.

Algorithm 1 Inverted Dynamic Routing (IDR). i and j are the indices of capsule heads in layer L and $L + 1$ respectively.

```

1: procedure IDR( $\mathbf{V}_{ij}^L, n_{iter}^L$ )                                ▶ given the votes and number of routing iterations
2:    $\mathbf{R}_{ij}^{\text{pre},L} \leftarrow 0$                                 ▶ initialize routing coefficients
3:   for  $n_{iter}^L$  iterations do
4:      $\mathbf{R}_{ij}^L \leftarrow \text{softmax}(\mathbf{R}_{ij}^{\text{pre},L})$                 ▶ Softmax among capsules in head  $i$ 
5:      $\tilde{\mathbf{A}}_{ij}^L \leftarrow \mathbf{R}_{ij}^L \odot \mathbf{V}_{ij}^L$                 ▶  $\odot$  is the Hadamard product
6:      $\mathbf{P}_j^{L+1} \leftarrow \text{squash}(\sum_i \sum_{xy} \tilde{\mathbf{A}}_{ij}^L)$         ▶  $\sum_i$  and  $\sum_{xy}$  are sums over heads and locations, respectively
7:      $\mathbf{R}_{ij}^{\text{pre},L} \leftarrow \mathbf{R}_{ij}^{\text{pre},L} + \mathbf{P}_j^{L+1} \cdot \mathbf{V}_{ij}^L$ 
8:   end for
9:    $A_{ij}^L \leftarrow \text{length}(\tilde{\mathbf{A}}_{ij}^L)$                         ▶ length computed by Eq. (1)
10:  return  $\mathbf{P}_j^{L+1}, A_{ij}^L$ 
11: end procedure

```

tion maps as the number of capsule heads per output class (see Fig. 3.2). Class-specific activation maps are the natural output of the proposed framework, unlike CNNs which require the use of additional modules, such as channel grouping, to cluster spatially-correlated patterns [32]. We utilize the activation maps to generate ROIs when an object is detected. This effectively yields a model capable of weakly-supervised localization which is trained end-to-end; we train on the images with categorical annotations and predict both the category and the *location* (i.e. mask or bounding box) for each test image. This framework is thus able to simultaneously generate multiple ROIs within the same image that are associated with different medical conditions.

The DECAPS Architecture: The overall architecture of the proposed DECAPS is shown in Fig. 1. It uses a ResNet [19] with three residual blocks as the base network which outputs 1024 feature maps, followed by a 1×1 convolutional layer with $A = 512$ channels and a ReLU nonlinearity. All the other layers are capsule layers starting with the primary capsule layer. The 4×4 pose matrix of each of the B primary capsule heads is a learned linear transformation of the output of all the lower-layer ReLUs centered at that location. The primary capsules are followed by two convolutional capsule layers with C and D capsule heads and kernels of size $K = 3$ and stride $s = 1$. We selected $B = C = D = 32$ capsule heads for the capsule layers, and used the proposed inverted dynamic routing (IDR) mechanism to rout the information between the capsules. The last layer of convolutional capsules is linked to the final dense capsule layer which has one capsule per output class. The Frobenius norm of the pose matrices of the output capsules are used to determine the predicted class.

Loss Function: We use spread loss to enforce the pose matrices of the top-level capsule j to have a large Frobenius norm if and only if the object of the corresponding class exists in the image [22]. The total spread loss is the sum of the losses for each output capsule as given by

$$L_{\text{spread}} = \sum_{j \neq t} L_j = \sum_{j \neq t} \max(0, m - (a_t - a_j))^2 \quad (2)$$

which directly maximizes the distance between the activation of the target class, a_t , and the activation of the other classes, a_i , using the margin m . In other words, it penalizes the model if the distance between the a_t and a_i is less than the margin m . We

initialize the margin to 0.2 and linearly increase it by adding 0.1 to its value every other epoch, which reduces the incidence of dead capsules in the hidden layers.

Activation-Guided Training (Peekaboo): To further promote DECAPS to focus on fine-grained details, we propose the Peekaboo strategy for capsule networks. Our strategy boosts the performance of DECAPS by forcing the network to look at all relevant parts for a given category, not just the most discriminative parts. It is inspired by the Hide and Seek [33] mechanism in which image patches are randomly dropped to encourage the network to look for other relevant parts. Peekaboo, however, uses the HAMs to guide the network’s attention process. For each training image, we randomly select an activation map A_{ij} in the range $[0, 1]$ to get the normalized HAM, $A_{ij}^* \in \mathbb{R}^{h_L \times w_L}$. We then enter a two step process: patch cropping, which extracts a fine-grained representation of the ROI to learn how to encode details, and patch dropping, which encourages the network to attend to multiple ROIs. In patch cropping, a mask $M_{ij}^{\text{crop}} \in \mathbb{R}^{h_L \times w_L}$ is obtained by setting all elements of A_{ij}^* which are less than a cropping threshold $\theta_c \in [0, 1]$ to 0, and 1 otherwise. We then find the smallest bounding box which covers the entire ROI, and crop it from the raw image (Fig. 1 (b)). It is then upsampled and fed into the network to generate a detailed description of the ROI. During the patch dropping procedure, M_{ij}^{drop} is used to remove the ROI from the raw image by using a dropping threshold $\theta_d \in [0, 1]$. The new patch-dropped image is then fed to the network for prediction. This encourages the network to train capsule heads to attend to multiple discriminative semantic patterns. Note that we flatten the output prediction matrices to form vectors (as shown in Fig. 1) whose magnitude represents the existence probability of the corresponding class.

At test time, we first input the whole image to obtain the coarse prediction matrices. ($\mathbf{p}_j^{\text{coarse}}$ for the j^{th} class) and the HAMs A_{ij} from all capsule heads. We then average all maps across the heads, crop and upsample the ROIs, and feed the regions to the network to obtain the fine-grained prediction vectors ($\mathbf{p}_j^{\text{fine}}$). The final prediction $\mathbf{p}_j^{\text{dist.}}$, referred to as distillation (Fig. 1 (c)), is the average of the $\mathbf{p}_j^{\text{coarse}}$ and $\mathbf{p}_j^{\text{fine}}$. The Peekaboo strategy can also be implicitly interpreted as an online data augmentation technique; it intelligently alters the input images (by either occluding or

passing up-sampled crops of regions of interest) to synthetically increase the number of samples passed through the network. We can therefore say that this feature helps to mitigate the effects of data scarcity.

3.3 Image-to-Image Conditional GAN Data Augmentation

We propose the use of the Image-to-Image (pix2pix) conditional GAN architecture as a means to augment our otherwise scarce COVID-19 CT scan dataset [34]. This architecture learns the mapping of both the input set (non-COVID-19 CT scans) and the output set (COVID-19 CT scans) by using a *learned* loss function, thus not requiring complex loss formulations. This allows us to learn the data distribution for each class; characteristics that make each class unique are effectively encoded and then sampled to create unique cases. We selected the pix2pix architecture due to its ability to retain stochasticity in the generated data, a feature that is lost in other conditional GAN architectures; this randomness is critical to the generation of unique cases that enrich our data. To accomplish this, we drew from the original implementation of pix2pix, substituting the Gaussian noise typically introduced in conditional GAN architectures with train-test dropout. Our generative model was also designed using the characteristic “U-net” architecture, where each layer i has a skip connection to layer $n - i$, with n being equal to the total number of layers [35]. This effectively prevents the network from minimizing the low-level information being passed, thus improving the overall quality of the output. Finally, we mitigate the negative effect the L1 loss has on image “crispness” by restricting the discriminative model’s attention to local image patches (i.e. the PatchGAN architecture). This effectively allows us to model high-level information while relying on the L1 loss to force low-level correctness, thus modeling the images as a Markov random field (i.e. simulating style loss). Additionally, we note that the PatchGAN architecture significantly reduces the number of trainable parameters and can therefore be applied to larger images (such as full resolution CT scans).

4 COVID-19 CT DATASET

In this paper, an open-source dataset prepared by Zhao *et al.* [36] is utilized to train and test the proposed model. It contains a total of 746 chest CT images, which are divided into two classes, namely COVID-19 and non-COVID-19. A pre-processed version of the dataset is available at <https://github.com/UCSD-AI4H/COVID-CT>.

The dataset was created by collecting images from papers related to COVID-19 published in medRxiv, bioRxiv, NEJM, JAMA, Lancet, and other impactful journals. Then, the images were classified according to the figure captions describing the clinical findings in the papers. All copyrights of the data belong to the authors and publishers of these papers. 349 CT images were labeled as COVID-19 and 397 CT images as non-COVID-19. The heights of these images vary; they range between 153 and 1853 pixels (average of 491 pixels), while their widths range between 124 and 1458 pixels (average of 383 pixels). We use nearly 85% of data for training/validation purposes and the rest for model testing. Specifically, the dataset is divided into a training set of 625 CT images (286 positives, 339 negatives) and a test set of 105 CT images (47 positives, 58 negatives).

5 EXPERIMENTAL RESULTS

Implementation details: All COVID-19 CT deep fakes were generated using the pix2pix architecture; images were re-scaled to 286×286 and then cropped following the PatchGAN mechanism to 256×256 . Each input (non-COVID-19) and output (COVID-19) was selected at random during the training procedure with replacement; this was done to maximize the stochasticity of the resulting augmented dataset. All hyperparameter selections were made to ensure that the generated output image was as sharp as possible for interpretation by experienced radiologists. We implemented the U-net 256 architecture as the image generator model using an Adam optimizer ($\beta_1 = 0.5$) and an initial learning rate of 0.0002 with a linear learning rate policy. All models were trained for 200 epochs with a batch size of 2; learning rate was kept constant for the first 100 epochs and then allowed to decay linearly to zero for the last 100 epochs. Finally, the L1 to discriminator loss weight ratio was 10 following the original implementation of pix2pix.

For DECAPS, the best performance was achieved using 3 routing iterations between the capsule layers. The Peekaboo hyperparameters were set as $\theta_c = 0.5$, and $\theta_d = 0.3$. The network is trained using the Adam optimizer with $\beta_1 = 0.5$, $\beta_2 = 0.999$, a learning rate of 10^{-4} (fixed for the duration of training) and a fixed batch size of 16. Input images are fed into the network with size 448×448 pixels yielding a 24×24 capsule map per head for the last convolutional capsule layer.

COVID-19 CT Deep Fakes: A total of 900 COVID-19 CT deep fakes were generated using the COVID-19 and non-COVID-19 samples; these images were later added to the original COVID-19 training set during the classification phase. We trained the pix2pix model until the resulting lung tissue information was sharp and qualitatively consistent with the training set images. To determine the semantic quality of the generated samples, we presented a small subset of the generated images to three experienced radiologists for evaluation. In Fig. 3, we note that the collaborating radiologists highlighted regions consistent with the manifestation of COVID-19 (red outlines) and areas that appear to resemble other diseases (blue outlines). All of the images in the subset were considered consistent with COVID-19 by at least one of the three radiologists. Given that the images in the non-COVID-19 set were not always control cases, some of the generated cases contained characteristics consistent with both COVID-19 and other diseases.

Architecture Comparison and Ablation Study: To compare the effectiveness of DECAPS to other traditional classification architectures, we trained the Inception-v3, DenseNet121, and ResNet50 architectures with the pix2pix augmented COVID-19 CT dataset. In Table 1, we note that DECAPS outperforms or matches all other architectures in all six statistical measures. The sharp increase in recall we observe indicates that the DECAPS architecture was more sensitive to COVID-19 cases when coupled with the Peekaboo training procedure. We then performed an ablation study to understand how inverted dynamic routing, pix2pix data augmentation, patch dropping, patch cropping, distillation, and their combination impacted the overall performance of the CapsNet classifier (as shown in Table 2). The sharp increase in performance attributed to the pix2pix image augmentation procedure implies that the technique partially

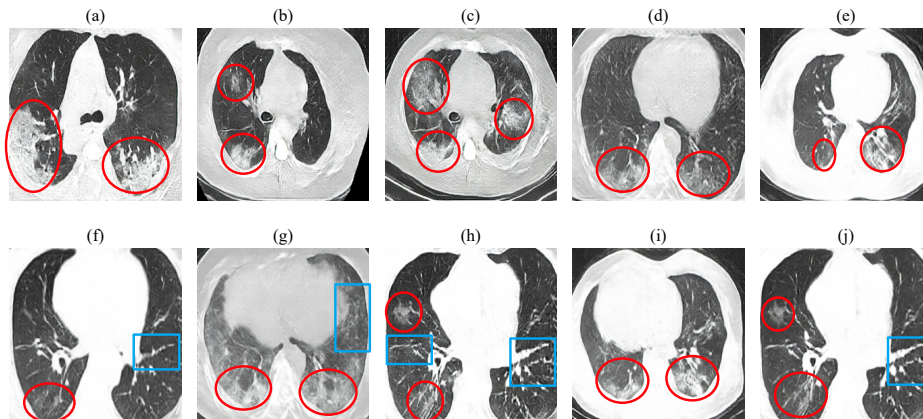


Figure 3: Qualitative assessment of COVID-19 CT pix2pix deep fakes. Red circles highlight characteristics consistent with the diagnosis of COVID-19, while blue squares correspond to features of non-COVID-19 diseases. We hypothesize that cases containing both kinds of features (“artificial co-morbidity” cases, (f), (g), (h), and (j)) help to enforce the unique characteristics of COVID-19.

Table 1: Prediction performance of models trained on the CT dataset. For each model, average (\pm std.) performance measure is reported over the best 5 trained model checkpoints.

	params. (M)	Precision	Recall	Specificity	Accuracy	F1-score	AUC
Inception-v3 [37]	25.1	0.844 (± 0.039)	0.740(± 0.119)	0.853(± 0.056)	0.819(± 0.028)	0.781(± 0.052)	0.894(± 0.029)
DenseNet121 [38]	7.0	0.815(± 0.063)	0.794(± 0.098)	0.839(± 0.074)	0.825(± 0.027)	0.801(± 0.038)	0.903(± 0.022)
ResNet50 [19]	23.5	0.7539(± 0.036)	0.849(± 0.121)	0.769(± 0.068)	0.808(± 0.039)	0.795(± 0.057)	0.880(± 0.030)
DECAPS	9.8	0.825(± 0.038)	0.854(± 0.079)	0.860 (± 0.074)	0.832(± 0.031)	0.837(± 0.031)	0.927(± 0.017)
DECAPS+Peekaboo	9.8	0.843 (± 0.024)	0.915 (± 0.057)	0.852(± 0.040)	0.876 (± 0.010)	0.871 (± 0.019)	0.961 (± 0.009)

addressed the data scarcity issues we mentioned in section 3. We note that each added component of the DECAPS architecture improves results, ultimately yielding a final AUC value of 96.1%.

Comparison to Thoracic Radiologists: To evaluate the diagnostic accuracies of the human radiologists and compare them to that of DECAPS, we created a high-quality test set (called Alpha-Test) by manually removing low-resolution images from the original test set and performing visual inspection to ensure that these images do not have artifacts such as radiographic annotations or texts that potentially bias the human radiologists. Our Alpha-Test eventually has 20 COVID positive cases and 32 non-COVID cases. Radiologists independently perform diagnosis on each image without knowing their ground truth information. We generate consensus diagnoses by computing the majority voting (e.g., diagnosis agreed by the majority of radiologists) to simulate the decision from a panel of radiologists. Fig. 4 shows the Receiver-Operating Characteristic (ROC) curve of

our network compared to radiologists’ performances. DECAPS + Peekaboo yields the best results with an area under curve (AUC) of 98%, and significantly outperforms all 3 radiologists. At a 15.64% false positive rate, the network achieves a 95% true positive rate compared to 85.11% true positive rate of the best radiologist. We also observe that the performance of the radiologist consensus is better than that of 2 radiologists and

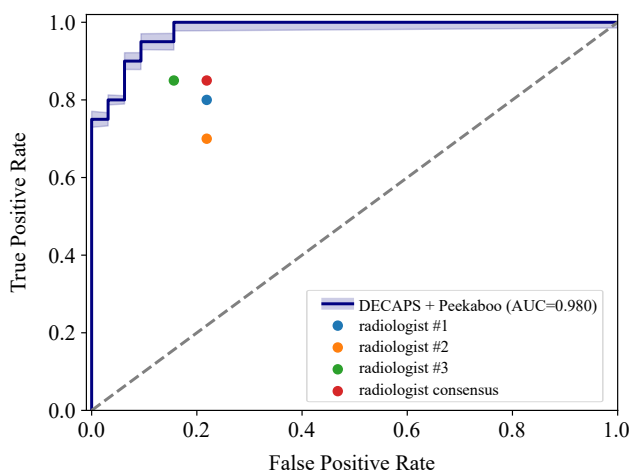


Figure 4: Performance comparison between the DECAPS+Peekaboo pix2pix augmented model performance and radiologist performance on the Alpha-Test set. We note that the performance of the radiologists (and their consensus performance) is significantly lower than the architecture.

Table 2: Mean AUC (\pm std.) over the top five models trained on the CT dataset to show the effect of each component and their combinations.

IDR	Pix2Pix	Patch Drop	Patch Crop	Distillation	AUC
					0.844(± 0.031)
✓					0.857(± 0.027)
✓	✓				0.874(± 0.041)
✓	✓				0.927(± 0.017)
✓	✓	✓			0.938(± 0.022)
✓	✓	✓	✓		0.944(± 0.017)
✓	✓	✓	✓		0.952(± 0.013)
✓	✓	✓	✓	✓	0.961(± 0.009)

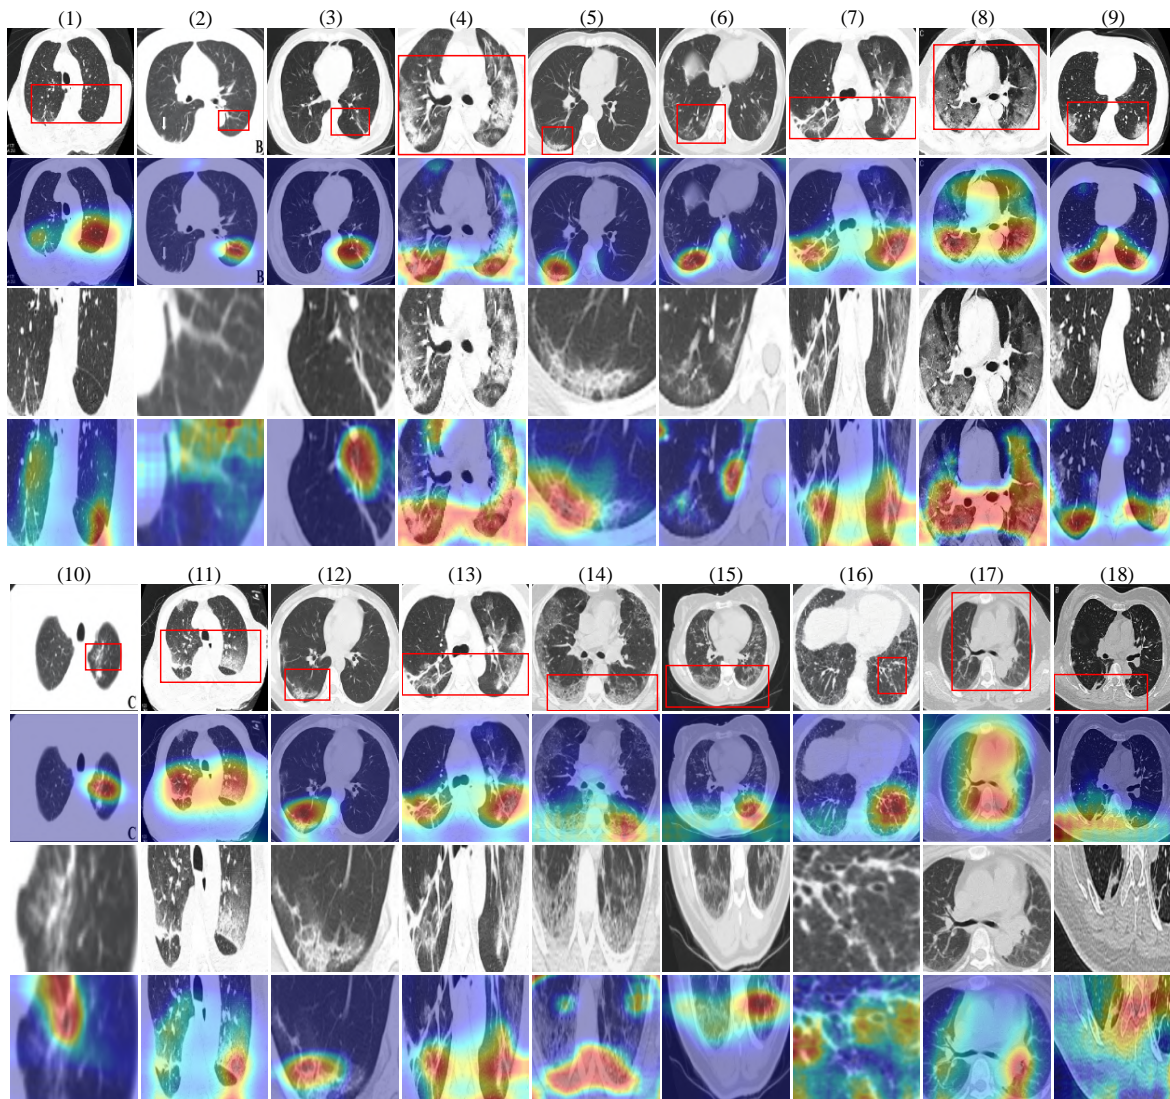


Figure 5: Samples describing the qualitative performance of the DECAPS architecture. Each column of 4 images presents the predicted bounding box on the input image, the coarse-grained class activation map, the cropped ROI, and the fine-grained class activation map, in that order. Samples (1)-(13) correspond to true positive cases, samples (14)-(16) correspond to true negative cases, and samples (17)-(18) correspond to failure cases.

worse than the best radiologist; at a 21.94% false positive rate, DECAPS + Peekaboo achieves a perfect true positive rate while the consensus achieves an 85.11% true positive rate.

DECAPS Visualization and Interpretation: We generated HAMs for all of the correctly classified COVID-19 cases (Fig. 5, cases 1-13) to visualize the information being used to generate the COVID-19 label. While most of the HAMs indicate that the correct features are being used, we note that sample 2 is classified using information that is not indicative of COVID-19 (a seemingly normal vascular branch). The samples correctly labeled as non-COVID-19 (Fig. 5, 14-16) show that the classifier was able to recognize disease chronicity, which is not indicative of acute infection. Finally, we note that both sample 17 and 18 from the misclassified sample group were classified using information that was irrelevant to the classification task.

6 DISCUSSION

COVID-19 has caused the world significant suffering; the widespread outbreak of the disease has effectively crippled the world healthcare and economic systems, demanding decisive action. While frightening, research shows that the rapid identification and isolation of COVID-19 cases can help to efficiently prevent spread [39]. While we would ideally use rRT-PCR to evaluate all potential cases, logistical issues have prevented this technique’s widespread adoption and use [40]. It would, therefore, be advantageous for hospitals to make use of resources that they already have; for instance, CT scans are used frequently to triage suspected COVID-19 patients. We were interested in developing a classification architecture that could improve upon the current state of the CT derived diagnostic performance in the context of COVID-19. To accomplish this, we combined a state-of-the-art conditional GAN (pix2pix) data augmentation

technique and a novel CapsNet architecture (DECAPS) that attempts to mimic the diagnostic procedures commonly used by radiologists.

Data scarcity is a limiting and difficult problem to solve in the context of AI; while classifiers can be improved and optimized to better understand the information they process, the data itself must adequately reflect some desired concept to be modeled. During a crisis like the COVID-19 pandemic, spending a significant amount of time to collect a well-vetted and large dataset is simply not feasible. This data scarcity is an important roadblock that must be rapidly surmounted to address the automation needs of the medical community in time. We were interested in leveraging the power of conditional GANs to help accomplish this; as has been shown in previous works, “deep fakes” could be used to increase the stochasticity of a dataset while retaining some underlying target concept. One can imagine that the generated samples resemble the examples produced by an educator; drawings from a medical textbook (while clearly fake) help to improve the understanding of the learning student. We found that the Image-to-Image (or pix2pix) architecture produced results that were consistent with what we observed in the COVID-19 CT dataset. To ensure that the generated samples captured the underlying concepts that distinguish COVID-19 scans from non-COVID-19 scans, we presented a small sample of the produced cases to three experienced radiologists (as shown in Fig. 3). While all three radiologists recognized that the samples were fake, each image was identified as having features consistent with COVID-19 by at least one of the physicians. Interestingly, they noted that several of the samples had features characteristic of other diseases (we refer to these cases as the “artificial co-morbidity” cases). For example, Fig. 3 (g) contains both tree-in-bud opacities in the left lung (blue circle, indicative of non-COVID-19) and ground glass and consolidation (red circles, indicative of COVID-19). Due to the positive effect these samples had on performance, we hypothesize that these cases helped to enforce the unique features of COVID-19 when training the DECAPS classifier.

The success of a classification architecture in radiography imaging tasks is dependent on the ability of the classifier to mimic the behavior of the radiologist. Being conscious of both the entire scan while focusing on a lesion (or lesions) of interest highlights the skill required for the accurate assessment of radiographic data. Most classification architectures are not designed to encourage this kind of behavior; typically, images are simply passed through the classification mechanism without explicitly being forced to consider details with respect to global information. The DECAPS architecture is designed to consider these relationships; it combines fine-grained and coarse information, generating classifications that are derived from the complex relationships between details and the big picture. Based on the generated performance metrics shown in Table 1, we see that the DECAPS architecture succeeds in capturing the characteristics of COVID-19. The sharp increase in recall attributed to the DECAPS architecture implies that the classifier was able to achieve a relatively profound understanding of the features unique to COVID-19. Our ablation study also shows how each unique component in the architecture contributes to this understanding, eventually reaching an AUC value of 96.1% (as shown in Table 2). The advantages of this architecture are most evident when compared to the performance of three experienced

radiologists; at a false positive rate of 15.64%, the top performing radiologist achieves a true positive rate of 85.11% while DECAPS+Peekaboo achieves a true positive rate of 95% (as shown in Fig. 4). We also note that the performance of the radiologist consensus is better than that of 2 of the participating radiologists and worse than the best radiologist; at a 21.94% false positive rate, DECAPS+Peekaboo achieves a perfect true positive rate while the consensus achieves an 85.11% true positive rate. Although further experimentation is required, these results imply that DECAPS could potentially match or exceed radiologist performance.

Inspection of all DECAPS HAMs by three experienced radiologists revealed that the classifier made use of information that was relevant to the classification of COVID-19 in most cases (Fig. 5, (3)-(13)). We note that the classifier was sensitive to peripheral ground-glass opacities which are considered a hallmark of COVID-19 in thoracic CT scans ((5), (6), (9), and (12) as examples). These results indicate that coupling the DECAPS architecture with the augmented dataset effectively addressed most of the underlying data scarcity issues. However, upon observation of sample (2), we note that the architecture used a seemingly normal vascular branch in the left lung to generate its correct COVID-19 classification. We hypothesize that the classifier mistook this feature as a linear band with adjacent ground-glass opacity, which would be indicative of late-stage COVID-19. Similarly, the architecture correctly detected a linear band with adjacent ground-glass opacity in sample (3). It is important to note that the classifier did not “cheat” by using the annotation in the right lung of sample (2) that highlighted the peripheral ground-glass opacity, proving that the architecture was seeking COVID-19 specific information. Another important case to consider is sample (1); we note that the classifier focused on a fissure in the left lung (a normal anatomical feature) while simultaneously considering the linear band with adjacent ground-glass opacities from the right lung. As mentioned earlier with sample (2), we believe that the classifier mistook the fissure as a feature of COVID-19. Interestingly, upon further examination, one of the radiologists changed their diagnosis based on the features being highlighted in the right lung; this highlights the importance of incorporating information from multiple regions of the scan to reach the correct diagnosis. This interaction between the physician and the AI system characterizes the advantages of having an AI-physician diagnostic team.

In samples (14)-(16), we observe cases where the classifier correctly identified traction bronchiectasis as a feature of a chronic fibrotic process, and thus non-COVID-19, despite the presence of ground-glass opacities. Finally, we see two failure cases in samples (17) and (18); while (17) is a non-COVID-19 sample with centrally located ground-glass opacities, sample 18 is a late-stage COVID-19 case mislabeled as non-COVID-19. We hypothesize that the dataset is not fully illustrative of late-stage COVID-19 features due to the nature of the disease and the low supply of descriptive data. The majority of COVID-19 publications have thus far focused on the early presentation of the disease, while patients improving or recovering from the disease usually do not require CT scans, making it difficult to acquire late-stage data.

At the time of writing this paper, healthcare professionals are still fighting against COVID-19 around the world. High-quality

data and ground truth annotations are not available in abundance, partly because physicians are too busy taking care of patients to provide clinical feedback. As a result, our study is limited to small training and test sets which potentially fail to represent the true patient population. Further studies on large-scale, multi-institutional, and well-curated datasets are important to confirm our findings. This being said, our results are promising; although our dataset was relatively scarce, we exceeded radiologist performance on a high-quality subset of our testing data. More importantly, our model was used by the radiologists to review and revise their decisions, symbolizing the importance of visualization. Given the severity and sense of urgency instilled by the COVID-19 pandemic, our results suggest that radiologists should utilize COVID-19 CT classifiers like DECAPS to review each case they encounter.

7 CONCLUSION

In this paper, we present Detail Oriented Capsule Networks (DECAPS), a unique capsule network architecture that is meant to mimic the complex behaviors of radiologists in the COVID-19 CT classification task. By coupling conditional generative adversarial networks (GANs) as a data augmentation technique with DECAPS, we effectively address the issue of data scarcity. Three well-trained and experienced radiologists then assessed the generated samples as a means of quality control. To further evaluate our model, we then compared its performance to the collaborating radiologists, showing that it outperformed all of them in a high-quality subset of our test set images. We then presented several activation maps to the radiologists for qualitative analysis of the classifier, and show that the information being used was relevant to the classification task. While further experimentation with a larger dataset is recommended, our results indicate that architectures like DECAPS could be used to assist radiologists in the CT mediated diagnosis of COVID-19.

REFERENCES

- [1] Search Results Web results COVID-19 RT-PCR Test - FDA. <https://www.fda.gov/media/136151/download>, 2020. Online; April 12 2020.
- [2] Tao Ai, Zhenlu Yang, Hongyan Hou, Chenao Zhan, Chong Chen, Wenzhi Lv, Qian Tao, Ziyong Sun, and Liming Xia. Correlation of chest ct and rt-pcr testing in coronavirus disease 2019 (covid-19) in china: a report of 1014 cases. *Radiology*, page 200642, 2020.
- [3] Yicheng Fang, Huangqi Zhang, Jicheng Xie, Minjie Lin, Lingjun Ying, Peipei Pang, and Wenbin Ji. Sensitivity of chest ct for covid-19: comparison to rt-pcr. *Radiology*, page 200432, 2020.
- [4] Michael Chung, Adam Bernheim, Xueyan Mei, Ning Zhang, Mingqian Huang, Xianjun Zeng, Jiufa Cui, Wenjian Xu, Yang Yang, Zahi A Fayad, et al. Ct imaging features of 2019 novel coronavirus (2019-ncov). *Radiology*, 295(1):202–207, 2020.
- [5] Fengxiang Song, Nannan Shi, Fei Shan, Zhiyong Zhang, Jie Shen, Hongzhou Lu, Yun Ling, Yebin Jiang, and Yuxin Shi. Emerging 2019 novel coronavirus (2019-ncov) pneumonia. *Radiology*, 295(1):210–217, 2020.
- [6] Yan Li and Liming Xia. Coronavirus disease 2019 (covid-19): Role of chest ct in diagnosis and management. *American Journal of Roentgenology*, pages 1–7, 2020.
- [7] Felix Chua, Darius Armstrong-James, Sujal R Desai, Joseph Barnett, Vasileios Kouranos, Onn Min Kon, Ricardo José, Rama Vancheeswaran, Michael R Loebinger, Joyce Wong, et al. The role of ct in case ascertainment and management of covid-19 pneumonia in the uk: insights from high-incidence regions. *The Lancet Respiratory Medicine*, 2020.
- [8] Geoffrey D Rubin, Christopher J Ryerson, Linda B Haramati, Nicola Sverzellati, Jeffrey P Kanne, Suhail Raoof, Neil W Schluger, Annalisa Volpi, Jae-Joon Yim, Ian BK Martin, et al. The role of chest imaging in patient management during the covid-19 pandemic: A multinational consensus statement from the fleischner society. *Chest*, 2020.
- [9] Mahmud Mossa-Basha, Carolyn C Meltzer, Danny C Kim, Michael J Tuite, K Pallav Kolli, and Bien Soo Tan. Radiology department preparedness for covid-19: Radiology scientific expert panel. *Radiology*, page 200988, 2020.
- [10] Linda Wang and Alexander Wong. Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest radiography images, 2020.
- [11] Lin Li, Lixin Qin, Zeguo Xu, Youbing Yin, Xin Wang, Bin Kong, Junjie Bai, Yi Lu, Zhenghan Fang, Qi Song, et al. Artificial intelligence distinguishes covid-19 from community acquired pneumonia on chest ct. *Radiology*, page 200905, 2020.
- [12] Shikha Chaganti, Abishek Balachandran, Guillaume Chabin, Stuart Cohen, Thomas Flohr, Bogdan Georgescu, Philippe Grenier, Sasa Grbic, Siqi Liu, François Melot, et al. Quantification of tomographic patterns associated with covid-19 from chest ct. *arXiv preprint arXiv:2004.01279*, 2020.
- [13] Xiaowei Xu, Xiangao Jiang, Chunlian Ma, Peng Du, Xukun Li, Shuangzhi Lv, Liang Yu, Yanfei Chen, Junwei Su, Guanqing Lang, et al. Deep learning system to screen coronavirus disease 2019 pneumonia. *arXiv preprint arXiv:2002.09334*, 2020.
- [14] Ophir Gozes, Maayan Frid-Adar, Hayit Greenspan, Patrick D Browning, Huangqi Zhang, Wenbin Ji, Adam Bernheim, and Eliot Siegel. Rapid ai development cycle for the coronavirus (covid-19) pandemic: Initial results for automated detection & patient monitoring using deep learning ct image analysis. *arXiv preprint arXiv:2003.05037*, 2020.
- [15] Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Katie Shpanskaya, et al. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225*, 2017.
- [16] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiography dataset with uncertainty labels

- and expert comparison. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 590–597, 2019.
- [17] Monika Grewal, Muktabh Mayank Srivastava, Pulkit Kumar, and Srikrishna Varadarajan. Radnet: Radiologist level accuracy using deep learning for hemorrhage detection in ct scans. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 281–284. IEEE, 2018.
- [18] Jeremy Fairbank, Amir Jamaludin, Meelis Lootus, Timor Kadir, Andrew Zisserman, Jill Urban, Michele Battié, Iain McCall, Genodisc Consortium, et al. Automation of reading of radiological features from magnetic resonance images (mri’s) of the lumbar spine without human intervention is comparable with an expert radiologist (issls prize-bioengineering science): 56. In *Spine Journal Meeting Abstracts*, page 56. LWW, 2017.
- [19] 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.
- [20] Joseph Paul Cohen, Paul Morrison, and Lan Dao. Covid-19 image data collection. *arXiv 2003.11597*, 2020. URL <https://github.com/ieee8023/covid-chestxray-dataset>.
- [21] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. Dynamic routing between capsules. In *Advances in Neural Information Processing Systems*, pages 3856–3866, 2017.
- [22] Geoffrey E Hinton, Sara Sabour, and Nicholas Frosst. Matrix capsules with EM routing. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=HJWLfgWRb>.
- [23] Kevin Duarte, Yogesh Rawat, and Mubarak Shah. Video-capsulenet: A simplified network for action detection. In *Advances in Neural Information Processing Systems*, pages 7610–7619, 2018.
- [24] Aryan Mobiny and Hien Van Nguyen. Fast capsnet for lung cancer screening. *arXiv preprint arXiv:1806.07416*, 2018.
- [25] Aryan Mobiny, Hengyang Lu, Hien V Nguyen, Badrinath Roysam, and Navin Varadarajan. Automated classification of apoptosis in phase contrast microscopy using capsule network. *IEEE transactions on medical imaging*, 39(1): 1–10, 2019.
- [26] David Rawlinson, Abdelrahman Ahmed, and Gideon Kowadlo. Sparse unsupervised capsules generalize better. *arXiv preprint arXiv:1804.06094*, 2018.
- [27] Dakai Jin, Ziyue Xu, Youbao Tang, Adam P Harrison, and Daniel J Mollura. Ct-realistic lung nodule simulation from 3d conditional generative adversarial networks for robust lung segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 732–740. Springer, 2018.
- [28] Changhee Han, Yoshiro Kitamura, Akira Kudo, Akimichi Ichinose, Leonardo Rundo, Yujiro Furukawa, Kazuki Umemoto, Yuanzhong Li, and Hideki Nakayama. Synthesizing diverse lung nodules wherever massively: 3d multi-conditional gan-based ct image augmentation for object detection. In *2019 International Conference on 3D Vision (3DV)*, pages 729–737. IEEE, 2019.
- [29] Amirata Ghorbani, Vivek Natarajan, David Coz, and Yuan Liu. Dermgan: Synthetic generation of clinical skin images with pathology. *arXiv preprint arXiv:1911.08716*, 2019.
- [30] Karim Ahmed and Lorenzo Torresani. Star-caps: Capsule networks with straight-through attentive routing. In *Advances in Neural Information Processing Systems*, pages 9098–9107, 2019.
- [31] Adam Kosioerek, Sara Sabour, Yee Whye Teh, and Geoffrey E Hinton. Stacked capsule autoencoders. In *Advances in Neural Information Processing Systems*, pages 15486–15496, 2019.
- [32] Heliang Zheng, Jianlong Fu, Tao Mei, and Jiebo Luo. Learning multi-attention convolutional neural network for fine-grained image recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 5209–5217, 2017.
- [33] Krishna Kumar Singh and Yong Jae Lee. Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In *2017 IEEE international conference on computer vision (ICCV)*, pages 3544–3553. IEEE, 2017.
- [34] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [35] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [36] Jinyu Zhao, Yichen Zhang, Xuehai He, and Pengtao Xie. Covid-ct-dataset: a ct scan dataset about covid-19. *arXiv preprint arXiv:2003.13865*, 2020.
- [37] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [38] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, volume 1, page 3, 2017.
- [39] Roy M Anderson, Hans Heesterbeek, Don Klinkenberg, and T Déirdre Hollingsworth. How will country-based mitigation measures influence the course of the covid-19 epidemic? *The Lancet*, 395(10228):931–934, 2020.
- [40] Jon Cohen and Kai Kupferschmidt. Countries test tactics in ‘war’ against covid-19, 2020.