

Highly Efficient Representation and Active Learning Framework for Imbalanced Data and Its Application to COVID-19 X-Ray Classification

Heng Hao · Sima Didari · Jae Oh Woo ·
Hankyu Moon · Patrick Bangert

Received: date / Accepted: date

Abstract We propose a highly data-efficient classification and active learning framework for classifying chest X-rays. It is based on (1) unsupervised representation learning of a Convolutional Neural Network and (2) the Gaussian Process method. The unsupervised representation learning employs self-supervision that does not require class labels, and the learned features are proven to achieve label-efficient classification. GP is a kernel-based Bayesian approach that also leads to data-efficient predictions with the added benefit of estimating each decision's uncertainty. Our novel framework combines these two elements in sequence to achieve highly data and label efficient classifications. Moreover, both elements are less sensitive to the prevalent and challenging class imbalance issue, thanks to the (1) feature learned without labels and (2) the Bayesian nature of GP. The GP-provided uncertainty estimates enable active learning by ranking samples based on the uncertainty and selectively labeling samples showing higher uncertainty. We apply this novel combination to the data-deficient and severely imbalanced case of COVID-19 chest X-ray classification. We demonstrate that only $\sim 10\%$ of the labeled data is needed to reach the accuracy from training all available labels. Its application to the COVID-19 data in a fully supervised classification scenario shows that our model, with a generic ResNet backbone, outperforms (COVID-19 case by 4%) the state-of-the-art model with a highly tuned architecture. Our model architecture and proposed framework are general and straightforward to apply to a broader class of datasets, with expected success.

Keywords Representation Learning · Active Learning · Gaussian Process · X-ray Imaging

Heng Hao · Sima Didari · Jae Oh Woo · Hankyu Moon · Patrick Bangert
Samsung SDS Research America

Heng Hao E-mail: h.heng@samsung.com · Sima Didari E-mail: s.didari@samsung.com · Jae Oh Woo E-mail: jaeoh.w@samsung.com · Hankyu Moon E-mail: hankyu.m@samsung.com · Patrick Banger E-mail: p.bangert@samsung.com

1 Introduction

Medical imaging is one of the major applications of computer vision technologies. The applications range from the most straightforward task of image classification (such as X-Ray, ultrasound, fundus) to image segmentation (anatomy or lesions), 3D imaging, and functional imaging (fMRI). Our focus in this paper is image classification of chest X-ray images from suspected COVID-19 patients. The dataset is described in detail in Section 4.1.

We demonstrate a classical application of computer vision-based image classification. There has been significant progress in the past decade both in terms of theoretical insights and classification accuracy due to the development and adoption of Deep Neural Network (DNN) models. The most well-established approach is the supervised training of Convolutional Neural Networks (CNN), which first identifies informative image features from multiple layers of convolutional filters fed to a small number of classification layers that produce category decisions. However, this popular approach typically requires large numbers of labeled images from each category to achieve an accuracy level useful for medical diagnosis. Data collection and labeling are often very costly. In some cases, it is not feasible to collect enough data for a quick automated diagnosis, as experienced in the time-critical cases of the COVID-19 pandemic. This leads to a highly imbalanced class distribution (“Normal” cases \gg “COVID-19” cases) that negatively impacts the decision accuracy. Given these practical challenges, we depart from the standard approach and propose a highly data-efficient methodology that can achieve the same level of accuracy using significantly fewer images and labels. It is based on CNN unsupervised representation learning hybrid with a Gaussian Process (GP) classifier. The GP-provided uncertainty estimates enable active learning by ranking unlabelled samples and selectively labeling samples showing higher uncertainty.

The CNN-GP hybrid model is trained in decoupled two steps. The first step is the **representation learning**, which refers to an unsupervised training step with the goal of extracting image features that are used in diverse downstream tasks. Among many different approaches for this challenging task, the discriminative approach focuses on designing an optimization problem to build a similarity or heuristic-based representation space. [20, 22, 30, 39]. The *contrastive loss* based learning as a discriminative approach has been applied very successfully and shows state-of-the-art performance in classification [3, 10, 12, 13, 24, 26, 34, 41, 54]. Saunshi et al. [45] also provides general theoretical insights of contrastive learning-based representation. Especially Chen et al. [13] confirms very high label-efficiency of the learned representation; it is able to achieve SoTA accuracy using as little as 1% training labels. These results have clear implications to data imbalance problem: the learned features (1) do not overfit dominant classes because the training does not use the class information (2) capture less dominant classes more efficiently. We empirically confirm this hypothesis in our experimental study 4.2.

The second step is the **Gaussian Process (GP)** classifier, a non-parametric Bayesian method, that can produce the prediction and its uncertainty in one shot. Many techniques have been developed to extract Bayesian uncertainty estimates from DNN [21, 23], and it was observed that the lower layers of a DNN for images may not benefit as much from a Bayesian treatment [33]. Similarly, GP has been used at the top of DNNs and has been applied to both classification and regression problems while producing uncertainty estimates [4, 8]. However, it was observed

that a DNN such as ResNet efficiently learns CIFAR10 with a test accuracy of more than 96%, while kernel methods such as GP can barely reach 80% [53]. It is well accepted [2] that better representation input to the kernel is crucial. The contrastive learning representation described above, which aims to aggregate points with similarity measured by distance, will be the ideal input to a GP with a distance-based RBF (radial basis function) kernel¹. This decoupled two-step training CNN-GP hybrid model framework, not only alleviates the inline training issue of GP but also extracts a predetermined lower dimension feature space for the GP, so that the GP classifier shows accuracy advantage over the Bayesian linear classifier, or the finite width non-linear neural network [7, 38]. The GP classifier also offers two special properties that make it well-suited for medical data analysis: (1) As a Bayesian method, it provides the prediction and its uncertainty in one-shot, which will greatly help medical decision making; (2) Compared to other non-Bayesian methods, it handles the issue of class imbalance more effectively [43], which is quite common in medical field. We verify in Section 4.3.




Fig. 1: Average classes test accuracy from active learning compared to random selection. The solid line and shaded area show the mean and standard deviation of five runs.

This process is repeated multiple times, with the train set gradually increasing in size until the model performance reaches a particular stopping criterion. Active learning will considerably facilitate real-world adoption of AI [1, 9, 28, 47, 50], especially in the medical field where data collection and labeling are quite expensive.

The **main contributions** of this paper are summarized as below:

- We provide a novel data-efficient framework leveraging the unsupervised contrastive representations and GP classifier. It leads to improvement over the state-of-the-art DNN model (COVID-Net) in COVID-19 chest X-ray classification.
- This combination of representation learning and GP provides an effective solution to the challenging issue of class imbalance, especially when data volume is small.

¹ The RBF kernel is $k(r) = \sigma^2 \exp(-r^2/2l^2)$, where r is the Euclidean distance between input vectors, l length-scaler and σ^2 variance-scaler are two hyper-parameters.

Active learning [18] is one of the most powerful techniques to improve data efficiency by saving labeling efforts. Its primary goal is to use the minimum amount of labels to reach maximum performance. We start with a small amount of labeled data (initial train set) to train the model. Then, we use an *acquisition function* (often based on the prediction uncertainty or entropy) to evaluate the unlabeled pool set, choose the most helpful pooling samples and ask the external *oracle* (generally a human) for the label. These newly labeled data are then added to the train set to train an updated

- Our approach also leads to an efficient CNN-GP active learning framework. Its application to the highly imbalanced COVID-19 chest X-ray imaging results in saving $\sim 90\%$ of the labels without any accuracy drop (see Figure 1).

2 Related Work on COVID-19 X-Ray Analysis

Since the pandemic, several studies have emerged focusing on gathering data and developing machine learning methods for COVID-19 detection from chest X-ray images. Wang et al. [52] introduced one of the first open-access benchmark datasets and tailored deep CNN networks to detect COVID-19 from chest X-ray images. The database and the tailored neural network model were referred to as COVIDx and COVID-Net, respectively. COVID-Net was designed by implementing a human-machine collaborative strategy based on specified performance metrics to detect COVID-19 cases. The final model architecture uses lightweight residual projection-expansion-projection-extension (PEPX) patterns, with a first-stage projection and expansion, a depth-wise representation, and a second-stage projection followed by an extension. Employing COVID-Net on COVIDx dataset, they reached a accuracy of 91%, 94%, and 95% for the detection of COVID-19 pneumonia, non-COVID pneumonia, and normal cases, respectively.

Multiple network architectures have been also proposed for the detection of COVID-19 by modifying existing deep architectures, such as CovXNet [36], truncated Inception Net [19], deep CNN combined with long short-term memory (LSTM) model [29], and a five layer CNN feature extractor model with a SVM classifier [40].

Besides efforts to develop specially tailored deep model architectures for COVID detection, several studies focused on increasing the accuracy of COVID-19 detection through the implementation of various transfer learning, fine-tuning, and data augmentation techniques for known architectures. Chowdhury et al. [14], Ucar and Korkmaz [51], Ozturk et al. [42], and Brunese et al. [6], implemented transfer learning and image augmentation to detect COVID-19 cases using pre-trained deep CNN architectures such as SqueezeNet, DenseNet201, VGG19 and DarkNet. Nour et al. [40] used Bayesian optimization for hyper-parameter tuning. Bressemer et al. [5] compared sixteen different known architectures of CNN on openly available COVID-19 Image Data Collection and showed that using deeper models does not always lead to higher accuracy in detecting COVID positive cases.

Our work does not focus on finding specialized architectures for the given task of X-ray imaging but instead on data representation and GP model that make a better use of given data and (often highly imbalanced) labels.

3 Methodology

We illustrate our training framework in Figure 2. The CNN-GP hybrid model is trained in two decoupled steps. (1) We first train a ResNet-50 representation generator with contrastive loss (SimCLR [13]) by feeding the train and pool set without any labels. As this is an unsupervised setting, we could feed the test set in training. However, to simulate the active learning scenario and perform a fair evaluation of the model performance, we exclude the test images from this step.

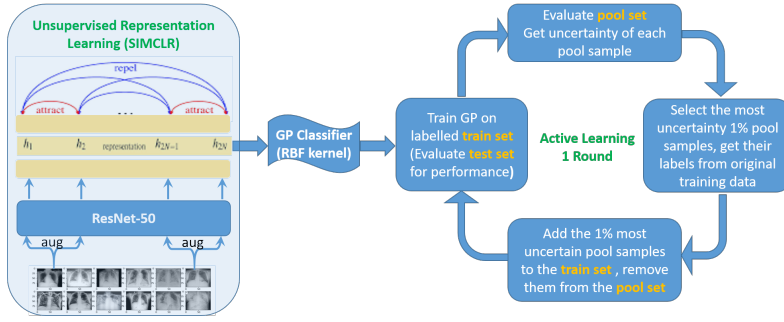


Fig. 2: Our active learning framework is illustrated. The representation generator is trained unsupervised with contrastive loss. The representations are used as inputs to the GP classifier. The GP classifier is trained in the active learning loop until the target performance is reached.

After the representation generator is trained, the train, pool, and test images are fed to the trained generator to produce representation vectors for the next step. (2) A multi-class GP classifier is trained using the training image representations with labels. The active learning cycle repeats step (2) while increasing the training sample size by labeling the samples selected via the acquisition function. Components of the framework are described in detail as following.

3.1 CNN-GP Hybrid Approach

We train the model by decoupling the representation and the classifier, following the practical guidance regarding the benefits of decoupling [32, 48]. Specifically, we adopt the *contrastive loss* to find a good representation space as proposed by Chen et al. [12, 13], followed by a separate GP classifier training.

3.1.1 SimCLR based Representation Learning

We start with a mini-batch with $N = 16$ image samples, image augmentation is applied twice to generate $2N$ samples. The image augmentation involves five operations - random crop, random flip, color distortion, Gaussian blur, and random gray-scale. To define a contrastive loss, we distinguish two types of augmented samples: positive pair and negative pair. Positive pairs are the ones augmented from the same image. For any other pairs, we consider them to be negative pairs regardless of the labels, which is unknown in the feature learning. The training maximizes the similarity of the positive pair, leveraging a contrastive loss. The contrastive loss between a pair of samples of i and j is given by

$$l_{i,j} = -\text{sim}(z_i, z_j)/\tau + \log \sum_{k=1}^{2N} \mathbf{1}_{[k \neq i]} \exp(\text{sim}(z_i, z_j)/\tau), \quad (1)$$

where $\text{sim}(\cdot, \cdot)$ is the cosine similarity between two vectors, and τ is a temperature hyperparameter. In SimCLR method [12, 13], the contrastive loss is evaluated at the projection head layer after the ResNet-50 backbone.

3.1.2 Gaussian Process

Once we calculate the representations (2,048 dimensions for each sample) of all the train set, we use them as the input to the GP classifier. To train the multi-class GP, we use the Sparse Variational GP (SVGP) [27] from GPflow package [37]. We choose the RBF kernel with 128 inducing points. We trained the model for 24 epochs using Adam optimizer with a learning rate of 0.001.

3.2 Active Learning Framework

We use the COVID-19 train set (13,942 samples, see Section 4.1) as the train (with labels) and pool sets (pretend we do not know the labels yet). We first randomly select 140 samples ($\sim 1\%$) as the initial train set for active learning. We train the GP classifier using the representations of the active learning train set with labels. We then use the trained GP to evaluate the test set and calculate the accuracy and confusion metrics to measure the GP classifier performance. Next, we use the same GP model to evaluate the pool data (the rest of the samples in the COVID-19 train set that is not in the current active learning train set) and calculate the prediction probabilities and uncertainties.

To select the most informative samples from the pool, various acquisition functions have been developed [21]. Here we considered the entropy, the uncertainty based on average class variance and the combination of both. Firstly, we compare the entropy of the pool samples. Entropy is calculated by $H(p) = -\sum_c p(c) \log p(c)$, where c is a class index [46]. Secondly, we compare the prediction uncertainties of the pool samples. For each sample, the GP classifier will provide the posterior variance of the prediction of each class. We calculate an average class variance, and consider the estimate to be the uncertainty of the pool sample. Lastly, considering both the entropy and the average class variance uncertainty, we get the sample’s entropy rank and the average class variance rank. We add the two rank numbers together as a combined rank. The $\sim 1\%$ pool samples with the largest entropy, average class variance or combined ranking are selected to be labeled and added to the active learning train set for the next round.

4 Experimental Results

4.1 COVID-19 Dataset




Fig. 3: From left to right, Normal, Pneumonia & COVID-19 chest X-rays.

COVIDx dataset is the largest open access benchmark dataset in terms of the number of COVID-19 positive patient cases to the best of our knowledge. At the time of this study, it consisted of 15,521 chest

X-ray images, of which 8,851 are “Normal”, 6,063 “Pneumonia” and 601 “COVID-19” pneumonia cases. A sample of these three types of X-ray images is shown in Figure 3. The dataset is the combination and modification of five different publicly available data repositories. These datasets are: (1) COVID-19 Image Data Collection [17], (2) COVID-19 Chest X-ray Dataset Initiative [16] (3) ActualMed COVID-19 Chest X-ray Dataset Initiative [15] (4) RSNA Pneumonia Detection Challenge Dataset, which is a collection of publicly available X-rays [44] , and (5) COVID-19 Radiography Database [31]. The dataset is highly imbalanced with significantly fewer COVID positive cases than other conditions. About 4% are COVID-19 positive cases. The train and test splits of the COVIDx dataset are depicted in Table 1. The class ratio of the three classes (“Normal”, “Pneumonia”, and “COVID-19”) for the train set is $\sim 16 : 11 : 1$ and for the test set is $\sim 9 : 6 : 1$.

	Normal	Pneumonia	COVID-19
Train	7,966	5,469	507
Test	885	594	100

Table 1: COVID-19 Dataset

Before feeding data to the representation generator, we pre-process the images by performing a 15% top crop, re-centering, and resizing to the original image size to delete the embedded textual information and enhance the region of interest [35, 49].

ded textual information and enhance the region of interest [35, 49].

4.2 Classification Results and Class Imbalance Handling

To compare with the state-of-art COVID-Net result [52], we train the CNN-GP hybrid model using all the COVID-19 training images (13,942 images). Note that the representation generator is still trained without any labels. The COVID-Net was trained using oversampling to balance the training classes. To do a fair comparison, we balance the training representations based on the corresponding labels before feeding to the GP classifier [11, 25]. In detail, we balance the representations by down-sampling “Normal” and “Pneumonia” classes and over-sampling “COVID-19” class so that the training size is kept constant while the difference in the sample sizes between classes is 1 or 2.

We report the total accuracy as **93.2%**, the average class accuracy as **93.6%**, and COVID-19 accuracy as **95%**, outperforming COVID-Net (see Table 4) by 4%. The confusion matrix is shown in Table 2. The resulting sensitivity is shown in Table 4. The normalized positive predictive value (PPV) is also laid out in Table 5 following the calculation of Wang et al. [52].

To validate the effectiveness of the unsupervised representation to data imbalance, we also train an NN based classifier (CNN-NN). We feed the same image representations to a single fully connected layer with softmax activation function, replacing the GP classifier. This NN classifier is trained 700 epochs with Adam optimizer. The confusion matrix is shown in Table 3. We report the total accuracy is **94.6%**, the average class accuracy is **93.9%**, and COVID-19 accuracy is **93%**. The classification results outperforms Covid-Net even with the simple ResNet-50 structure. This result implies our representation is rich enough to perform a linear classification task, and it significantly reduces the complexity of the network.

Actual \ Prediction	Prediction		
	Normal	Pneumonia	COVID-19
Normal	831 (93.9%)	51 (5.8%)	3 (0.3%)
Pneumonia	42 (7.1%)	546 (91.9%)	6 (1.0%)
COVID-19	3 (3.0%)	2 (2.0%)	95 (95.0%)

Table 2: Confusion matrix from CNN-GP classifier

Actual \ Prediction	Prediction		
	Normal	Pneumonia	COVID-19
Normal	849 (95.9%)	33 (3.7%)	3 (0.3%)
Pneumonia	37 (6.2%)	551 (92.8%)	6 (1.0%)
COVID-19	3 (3.0%)	4 (4.0%)	93 (93.0%)

Table 3: Confusion matrix from the NN (softmax) classifier

Architecture	Normal	Pneumonia	COVID-19	Average
VGG-19	98.0%	90.0%	58.7%	82.2%
ResNet-50	97.0%	92.0%	83.0%	90.7%
COVID-Net	95.0%	94.0%	91.0%	93.3%
CNN-NN (ours)	95.9%	92.8%	93.0%	93.9%
CNN-GP (ours)	93.9%	91.9%	95.0%	93.6%

Table 4: Sensitivity for each class. Best result is in **bold**. Other architecture results came from Table 2 of Wang et al. [52].

Architecture	Normal	Pneumonia	COVID-19
VGG-19	83.1%	75%	98.4%
ResNet-50	88.2%	86.8%	98.8%
COVID-Net	90.5%	91.3%	98.9%
CNN-NN (ours)	91.4%	92.1%	98.9%
CNN-GP (ours)	90.4%	92.0%	99.0%

Table 5: Normalized positive predictive value (PPV) for each class. Best result is in **bold**. Other architecture results are from Wang et al. [52] Table 3.

Both the two stage trained hybrid classifiers CNN-GP and CNN-NN outperform and show more balanced accuracy over the three classes compared to the other three fully-supervised models: VGG-19, ResNet-50, and COVID-Net.

4.3 Benefit of GP over NN for Imbalanced Data

To show the benefit of GP classifier over NN for imbalanced data, we compare the CNN-GP classifier with the CNN-NN classifier with the **same** fewer samples randomly selected from the training dataset. The CNN-NN classifier is trained in the same manner as laid out in the previous section. The CNN-GP reveals better data-efficiency compared to the CNN-NN classifier (see Figure 1 and Table 6). Because GP classifiers fit each class separately, as long as each class have

Train	CNN-NN	CNN-GP
10%	65.0%±10.7%	64.8%±9.1%
20%	72.8%±5.3%	77.2%±5.0%
30%	75.4%±1.5%	81.2%±1.8%
40%	78.4%±5.4%	82.6%±2.3%
50%	77.4%±1.5%	83.8%±2.4%

Table 6: The mean and standard deviation of the test **Covid-19 accuracy** for random selected train samples for five runs. Note, no active learning is performed in this experiment.

enough sample to show its own distribution, it is less affected by the imbalance rate between classes. We confirm when we have enough training samples ($> 10\%$) for the classifiers, CNN-GP classifier shows higher test COVID-19 accuracy (by $\sim 4\% - 6\%$) compared to the CNN-NN cases. Based on GP’s more robust behavior and less fluctuations of accuracy, we select the CNN-GP model as part of our active learning framework.

4.4 Active Learning Iteration

The active learning accuracy from the three acquisition functions (entropy, uncertainty, and combined ranking) are compared to the random selection in Figure 4. To check the consistency of our results, we repeat multiple (five) active learning runs. The results shown in the figure are the means of the five independent active learning runs. With the unsupervised representation learning followed by a GP classifier, only $\sim 10\%$ of the training data needs to be labeled to achieve the same accuracy as if all the labeled training data is used.

In Figure 4, the different line colors and styles illustrate active learning results from different acquisition functions but from the same CNN-GP hybrid model. The three different acquisition functions have similar performance (red dash, green solid and purple dash-dot lines) and outperforms the random selection (blue dotted line). We observe that, especially when the sample size is small ($< 20\%$), the training data selected by these three acquisition functions accelerates the model to reach significantly higher test accuracy. The remaining 90% of the data offer no new information to the classification model and can be auto-labeled by the CNN-GP hybrid model, saving considerable labeling cost.

We observe that if we balance the class representations before training the GP classifier, it will slightly help increasing the test accuracy. Thus, at each active learning round, we down-sampling “Normal”, “Pneumonia” classes and over-sampling “COVID-19” class similar to Section 4.2. We also run five independent iterations for this setup. The results are shown in bottom row of Figure 4.

5 Conclusion

We introduced a data-efficient CNN-GP hybrid model and showed that our approach achieves the state-of-the-art accuracy for COVID-19 chest X-ray classification. Moreover, our novel combination of representation learning and GP is

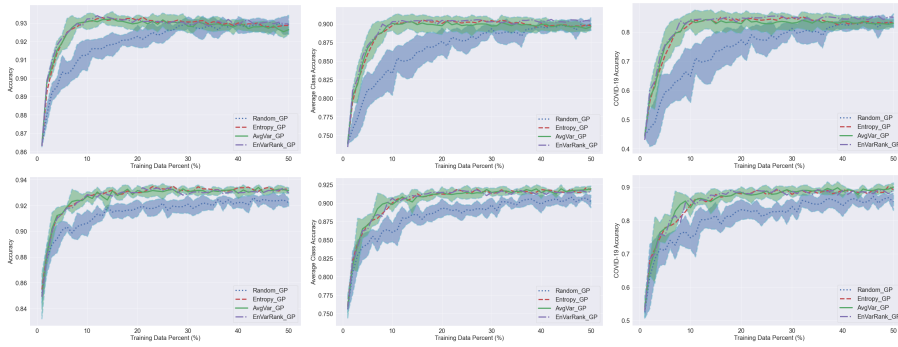


Fig. 4: The accuracy(left), average class accuracy (middle) and COVID-19 accuracy(right) for five active learning runs with different acquisition functions compared to the random pick. The lines show the mean and shaded areas show the standard deviation. The results without over-sampling the representations before feed in to the classifiers are shown in the top row, while the results from over-sampling of representations before feed in to the classifiers are shown in the bottom row.

confirmed to be an effective solution for the issue of class imbalance, especially when facing data scarcity as in COVID-19 case. Our approach also enables an efficient CNN-GP active learning with its application to the highly imbalanced COVID-19 chest X-ray imaging, leading to saving $\sim 90\%$ of the labeling time and cost.

Further improvement of the proposed framework is attainable through improved unsupervised representation learning and implementation of better acquisition functions with stronger exploration and exploitation characteristics. The aforementioned directions will be the focus of our future studies.

References

1. Adomavicius G, Tuzhilin A (2005) Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge & Data Engineering*
2. Allen-Zhu Z, Li Y (2019) What can resnet learn efficiently, going beyond kernels? *ArXiv abs/1905.10337*
3. Bachman P, Hjelm RD, Buchwalter W (2019) Learning representations by maximizing mutual information across views. In: *Advances in Neural Information Processing Systems*, pp 15535–15545
4. Bradshaw J, Matthews A, Ghahramani Z (2017) Adversarial examples, uncertainty, and transfer testing robustness in gaussian process hybrid deep networks. *arXiv preprint arXiv:170702476*
5. Bressen KK, Adams L, Erxleben C, Hamm B, Niehues S, Vahldiek J (2020) Comparing different deep learning architectures for classification of chest radiographs. *Scientific Reports* 10(13590)

6. Brunese L, Mercaldo F, Reginelli A, Santone A (2020) Explainable deep learning for pulmonary disease and coronavirus covid-19 detection from x-rays. *Computer Methods and Programs in Biomedicine* 196:105608, DOI 10.1016/j.cmpb.2020.105608
7. Bui T, Hernandez-Lobato D, Hernandez-Lobato J, Li Y, Turner R (2016) Deep gaussian processes for regression using approximate expectation propagation. In: *Proceedings of The 33rd International Conference on Machine Learning*, PMLR, New York, New York, USA, *Proceedings of Machine Learning Research*, vol 48, pp 1472–1481
8. Calandra R, Peters J, Rasmussen CE, Deisenroth MP (2016) Manifold gaussian processes for regression. In: *2016 International Joint Conference on Neural Networks (IJCNN)*, pp 3338–3345, DOI 10.1109/IJCNN.2016.7727626
9. Calinon S, Guenter F, Billard A (2007) On learning, representing, and generalizing a task in a humanoid robot. *IEEE Transactions on Systems, Man and Cybernetics, Part B (Cybernetics)* 37(2):286–298
10. Caron M, Misra I, Mairal J, Goyal P, Bojanowski P, Joulin A (2020) Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems* 33
11. Chawla NV (2009) Data mining for imbalanced datasets: An overview. *Data mining and knowledge discovery handbook* pp 875–886
12. Chen T, Kornblith S, Norouzi M, Hinton G (2020) A simple framework for contrastive learning of visual representations. *Proceedings of the International Conference on Machine Learning (ICML)*
13. Chen T, Kornblith S, Swersky K, Norouzi M, Hinton GE (2020) Big self-supervised models are strong semi-supervised learners. *Advances in Neural Information Processing Systems* 33
14. Chowdhury ME, Rahman T, Khandakar A, Mazhar R, Kadir MA, Mahbub ZB, Islam KR, Khan MS, Iqbal A, Al-Emadi N, et al. (2020) Can ai help in screening viral and covid-19 pneumonia? *IEEE Access* 8:132665–132676, DOI 10.1109/ACCESS.2020.3010287
15. Chung A (2020) Actualmed covid-19 chest x-ray data initiative. <https://github.com/agchung/Actualmed-COVID-chestxray-dataset>
16. Chung A (2020) Figure1-covid-chestxray-dataset. <https://github.com/agchung/Figure1-COVID-chestxray-dataset>
17. Cohen JP, Morrison P, Dao L (2020) Covid-19 image data collection. 2003. 11597
18. Cohn DA, Ghahramani Z, Jordan MI (1996) Active learning with statistical models. *Journal of Artificial Intelligence Research* 4:129–145
19. Dipayan D, C SK, Umapada P (2020) Truncated inception net: Covid-19 outbreak screening using chest x-rays. *Physical and Engineering Sciences in Medicine*
20. Doersch C, Gupta A, Efros AA (2015) Unsupervised visual representation learning by context prediction. In: *Proceedings of the IEEE international conference on computer vision*, pp 1422–1430
21. Gal Y, Ghahramani Z (2016) Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: *Proceedings of The 33rd International Conference on Machine Learning*, pp 1050–1059
22. Gidaris S, Singh P, Komodakis N (2018) Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:180307728*

23. Graves A (2011) Practical variational inference for neural networks. *Advances in Neural Information Processing System* 24:2348–2356
24. Grill JB, Strub F, Althé F, Tallec C, Richemond P, Buchatskaya E, Doersch C, Avila Pires B, Guo Z, Gheshlaghi Azar M, et al. (2020) Bootstrap your own latent—a new approach to self-supervised learning. *Advances in Neural Information Processing Systems* 33
25. He H, Garcia EA (2009) Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering* 21(9):1263–1284
26. He K, Fan H, Wu Y, Xie S, Girshick R (2020) Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp 9729–9738
27. Hensman J, Fusi N, Lawrence ND (2013) Gaussian processes for big data. In: *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence (UAI2013)*
28. Hoi SC, Jin R, Zhu J, Lyu MR (2016) Batch mode active learning and its application to medical image classification. In: *Proceedings of the 23rd international conference on Machine Learning*, pp 1492–1501
29. Islam MZ, Islam MM, Asraf A (2020) A combined deep cnn-lstm network for the detection of novel coronavirus (covid-19) using x-ray images. *Informatics in medicine unlocked*
30. Isola P, Zhu JY, Zhou T, Efros AA (2017) Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 1125–1134
31. Kaggle (2019) Radiological society of north america. covid-19 radiography database. <https://www.kaggle.com/tawsifurrahman/covid19-radiography-database>
32. Kang B, Xie S, Rohrbach M, Yan Z, Gordo A, Feng J, Kalantidis Y (2019) Decoupling representation and classifier for long-tailed recognition. In: *International Conference on Learning Representations*
33. Kendall A, Badrinarayanan V, Cipolla R (2017) Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. In: Tae-Kyun Kim GB Stefanos Zafeiriou, Mikolajczyk K (eds) *Proceedings of the British Machine Vision Conference (BMVC)*, BMVA Press, pp 57.1–57.12, DOI 10.5244/C.31.57, URL <https://dx.doi.org/10.5244/C.31.57>
34. Li J, Zhou P, Xiong C, Socher R, Hoi SC (2020) Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:200504966*
35. Maguolo G, Nanni L (2020) A critic evaluation of methods for covid-19 automatic detection from x-ray images. *arXiv preprint arXiv:200412823*
36. Mahmud T, Rahman MA, Fattah SA (2020) Covxnet: A multi-dilation convolutional neural network for automatic covid-19 and other pneumonia detection from chest x-ray images with transferable multi-receptive feature optimization. *Computers in Biology and Medicine* 122:103869, DOI 10.1016/j.combiomed.2020.103869
37. Matthews AGdG, van der Wilk M, Nickson T, Fujii K, Boukouvalas A, León-Villagrà P, Ghahramani Z, Hensman J (2017) GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research* 18(40):1–6, URL <http://jmlr.org/papers/v18/16-537.html>

38. Neal RM (1996) Bayesian Learning for Neural Networks. Springer-Verlag, Berlin, Heidelberg
39. Noroozi M, Favaro P (2016) Unsupervised learning of visual representations by solving jigsaw puzzles. In: European conference on computer vision, Springer, pp 69–84
40. Nour M, Cömert Z, Polat K (2020) A novel medical diagnosis model for covid-19 infection detection based on deep features and bayesian optimization. Applied Soft Computing 97:106580, DOI <https://doi.org/10.1016/j.asoc.2020.106580>, URL <http://www.sciencedirect.com/science/article/pii/S1568494620305184>
41. Oord Avd, Li Y, Vinyals O (2018) Representation learning with contrastive predictive coding. arXiv preprint arXiv:180703748
42. Ozturk T, Talo M, Yildirim EA, Baloglu UB, Yildirim O, Rajendra Acharya U (2020) Automated detection of covid-19 cases using deep neural networks with x-ray images. Computers in Biology and Medicine 121:103792, DOI <https://doi.org/10.1016/j.compbiomed.2020.103792>, URL <http://www.sciencedirect.com/science/article/pii/S0010482520301621>
43. Rosevear D, de Waal A (2017) Gaussian processes applied to class-imbalanced datasets
44. RSNA (2019) Radiological society of north america. rsna pneumonia detection challenge. <https://www.kaggle.com/c/rsna-pneumonia-detection-challenge/data>
45. Saunshi N, Plevrakis O, Arora S, Khodak M, Khandeparkar H (2019) A theoretical analysis of contrastive unsupervised representation learning. In: International Conference on Machine Learning, pp 5628–5637
46. Shannon CE (1948) A mathematical theory of communication. The Bell System Technical Journal 27:379–423, 623–656, URL <http://cm.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf>
47. Siddhant A, Lipton ZC (2018) Deep bayesian active learning for natural language processing: Results of a large-scale empirical study. ArXiv abs/1808.05697
48. Tang K, Huang J, Zhang H (2020) Long-tailed classification by keeping the good and removing the bad momentum causal effect. Advances in Neural Information Processing Systems 33
49. Tartaglione E, Barbano CA, Berzovini C, Calandri M, Grangetto M (2020) Unveiling covid-19 from chest x-ray with deep learning: a hurdles race with small data. arXiv preprint arXiv:200405405
50. Tong S (2001) Active learning: theory and applications. PhD thesis, Stanford University
51. Ucar F, Korkmaz D (2020) Covidiagnosis-net: Deep bayes-squeezenet based diagnosis of the coronavirus disease 2019 (covid-19) from x-ray images. Medical Hypotheses 140:109761, DOI <https://doi.org/10.1016/j.mehy.2020.109761>, URL <http://www.sciencedirect.com/science/article/pii/S0306987720307702>
52. Wang L, Lin ZQ, Wong A (2020) Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. Scientific Reports 10(1):1–12
53. Wilson AG, Hu Z, Salakhutdinov RR, Xing EP (2016) Stochastic variational deep kernel learning. In: Advances in Neural Information Processing Systems,

-
- pp 2586—2594
54. Zhuang C, Zhai AL, Yamins D (2019) Local aggregation for unsupervised learning of visual embeddings. In: Proceedings of the IEEE International Conference on Computer Vision, pp 6002–6012