SCIENTIFIC PAPER

COVID-19 AI PROJECT: AI Biokinetic Technologies' proprietary Horus-1 artificial intelligence automatically detects the novel COVID-19 using radiographic images through transfer learning and convolution neural networks as point of care application on a mobile device.

Version: 17-Apr-2020 by AI Biokinetic Technologies

Authors: Wayne Croft, PhD MMS; Kordel K. France; Ian Dawe, APRN FNP-C; Zach Newman; Jon Ward.

Abstract

In this proof of concept study, a dataset of radiographic images (X-ray, CT) from patients with, confirmed COVID-19 infections, Severe Acute Respiratory Syndrome (SARS), Common Bacterial Pneumonia (CBP), Emphysema, Pulmonary Nodules (PN), Atelectasis, Pneumothorax (PTX), and normal healthy subjects were evaluated by an artificial intelligence (AI) algorithm we denote as Horus-1. The procedure that was used to perform the investigation is a state-of-the art convolutional AI neural network compressed to a hand-held Apple mobile device. The primary objective of this study was to gauge if Horus-1 was able to detect COVID-19 lung scans intermingled with a variety of other lung diseases. The X-ray and CT radiographic images were obtained from public medical repositories located on the National Institutes of Health GitHub domain. Over 100,000 radiographic images were analyzed by Horus-1 to evaluate the efficacy of its ability to distinguish COVID-19 from other respiratory conditions. A total of 3,978 COVID-19 images were intermingled amongst images from the aforementioned various lung diseases. A specific procedure utilized to mitigate false negatives was hierarchical neural networks. As Horus-1 continued with its machine learning it achieved an average of 96% accuracy in differentiating COVID-19 scans from other X-ray and CT scans. Horus-1's precision and recall have consistently achieved a range of 94% to 98%. In a practical application, Horus-1 achieves this accuracy consistently in an average of 3.5 seconds. The results of the investigation suggest that Horus-1's deep learning of X-ray and CT images can identify specific disease patterns detectable as signatures or biomarkers related to COVID-19 infection in a practical real-world setting, and its architecture is? one that is successful in minimizing false-negative readings.

Key words: Artificial Intelligence, Automatic Detections, Computer Applications, COVID-19, Deep Learning, Hierarchical Learning, Transfer Learning, Radiographic Images, Edge Computing, Point of Care Application.

Introduction

Artificial Intelligence (AI) is one of the fastest growing fields of informatics significantly impacting the radiological sciences in a similar way the computer has impacted the laboratory science. Using machine learning to interpret all forms of radiographic and magnetic resonance images may be transformative in diagnosing diseases by identifying the signature or biomarker at the granular level that a human cannot do.

AI represents the capacity of machines to mimic the cognitive functions of humans (in this context, learning and problem solving). It can be further subdivided into artificial bio-intelligence, where a computer can perform a very specific task as well as or better than a human. AI can be understood as a set of tools and algorithms that make software "smarter" to the extent that an outside observer could compare the output to that generated by a human. It operates similarly to the way a highly intelligent human brain functions doing regular tasks like common-sense reasoning, forming an opinion, or even social behavior.[2]

The term machine learning implies the situation in which AI is learning if it improves its performance on future tasks after making observations about the data fed to it.[3] Machine learning algorithms evolve with increasing exposure to data; they are not based exclusively on rules, but improve with experience, learning to give specific answers by evaluating large amounts of data.[4.] The learning can be unsupervised, reinforced, supervised, and semi-supervised. In unsupervised learning, the AI learns patterns in the input even though no explicit feedback is supplied. As a general rule, the "deeper" the network (more layers) and the more rounds of training, the better the performance of the network. Where deep learning differs differentiates itself is through several network layers and generally many CPU or GPU cores working in tandem to train an extremely large network.

Several machine learning procedures were used to evolve an algorithm that the authors denote as *Horus-1.* These procedures are expounded upon throughout the following pages. The final algorithm is a series of seven convolutional neural networks compressed small enough to fit onto an Apple iPhone and iPad. The algorithm eventually proved successful in distinguishing the novel COVID-19 virus from 6 other different conditions in radiographic images of the human chest.

COVID-19 is a mild to severe respiratory illness that is caused by a coronavirus (Severe acute respiratory syndrome coronavirus 2 of the genus Beta-coronavirus). It is transmitted chiefly by contact with infectious material (such as respiratory droplets), and is characterized especially by fever, cough, and shortness of breath and may progress to pneumonia and respiratory failure.[4] Severe symptoms associated with massive alveolar damage results in death in approximately 3% of population.

It is estimated that there are three strains of the virus infecting patients and all three strains of the virus are infecting people over various parts of the globe. Experts from Cambridge University were able to map the genetic history of COVID-19 from December, when the spread was first recorded, to March, discovering that the virus has three distinct strains, all of which are closely related to each other. Analysis of the genetic history of SARS-CoV-2 or the novel coronavirus COVID-19 shows that Type A, is the strain that came from bats and transferred to humans from pangolins. Type B, which is derived from Type A coronavirus via two mutations, is more prevalent in the country, specifically in ground zero, Wuhan China. Type B of the mutation is affecting the West Coast of the USA and Australia. Type C coronavirus strain is also common in some places in Europe as well as in Asia, specifically in Singapore. This strain is the "daughter" of Type B and is just one mutation different from it.[5.] Recent investigative research suggests that Type A COVID-19 may have come from the Wuhan biolab and not the nearby wet-market as initially claimed by the Chinese Communist Party.

Regardless of the strain of the COVID-19 virus, the biomarkers or signature that it leaves in the lungs are similar. Thereby, Horus-1 appears to be able to accurately identify the biomarker on an average of 96% of the time. However, further research is necessary to determine when Horus-1 can identify the biomarker in the lungs after the patient has been infected. Preliminary research seems to show that Horus-1 can detect the biomarker as early as 2 days post infection.

Early and accurate detection of an infected COVID-19 patient is vital in order to mitigate the spread of the virus. Once a patient is suspected and diagnosed then quarantine is recommended due to a current lack of scientifically proven vaccinations and/or medication treatments being available. Clinical applications of case study off label use of the drug combination of Hydroxychloroquine and Zithromax appears to reduce the symptoms of COVID-19 and shortens the recovery time.

Utilizing Horus-1 to automatically identify the COVID-19 biomarkers in infected patients early in the process may prove beneficial for clinicians to quarantine and treat patients with better outcome measures. Head-to-head comparison studies of Horus-1 versus standard laboratory diagnostics needs to occur.

Several similar efforts that have been proposed for using AI in performing diagnostics on radiology images. However, these approaches seem to either utilize an internet connection to perform the computation on the cloud or use a stationed desktop computer to perform the diagnostics. What makes Horus-1 particularly unique compared to other AI technologies is that it can be compressed and utilized through an app on a mobile computing device such as an iPhone, iPad, or Google Pixel without the need of an internet connection and making it a point of care application. Additionally, the app contains options to allow the user to either "capture an image" or "import an image" of the X-ray. By compressing Horus-1 to a

handheld device, one enables access to the technology more readily to medical professionals, security professionals, and the patient to identify different COVID-19 biomarkers in real time. Not requiring dependence on cloud processing also opens up access of the screening technology to patients and individuals in rural areas. Furthermore, in the medical field, false negatives for critical diseases such as COVID-19 pose extremely dire consequences. In order to address this, AI Biokinetic Technologies constructed a specific neural network architecture to substantially mitigate false negative results.

We see this technology being utilized in hospitals, as well as border checkpoint areas of countries, state lines, or even airports and intercommunity public transit hubs, essentially anywhere a portable X-ray machine can be placed. In order to do this, AI needs to be made more readily accessible outside of supercomputers and cloud networks. With Horus-1, this technology is available in the palm of the user's hand with the output biomarker results within 3.5 seconds once the radiographic image is uploaded. Furthermore, since the computation is done on the mobile device itself, proper security is ensured over the patient's data since there is no uploading of an image over a signal connection, thereby making it more protective of patient's personal health information.

Methods

Dataset of study

For the purpose of the experiments completed in this proof of concept study, several sources of X-ray and CT scans were accessed. Datasets of radiographic images (X-ray, CT) from patients with, confirmed COVID-19 infections, Severe Acute Respiratory Syndrome (SARS), Common Bacterial Pneumonia (CBP), Emphysema, Pulmonary Nodules (PN), Atelectasis, Pneumothorax (PTX), and normal healthy subjects were fed into Horus-1's neural network and evaluated. The radiographic images were obtained from three open-sourced datasets including the National Institutes of Health via Kaggle[7] and two GitHub dataset forums[8,9]. In the case of the COVID-19 data, the population size was lower than preferred for the accuracy desired, so the dataset was further augmented using standard augmentation techniques.

Preliminary Results

Due to the nature of the problem being solved, the analysis team knew the architecture would have to be moderately complex with many hours of deep learning incorporated. However, in order to understand how recall and precision behaved in the network as it scaled, the team began with a simple, shallow network and analyzed results as complexity increased. We also began with a binary classification network distinguishing only two conditions, those images containing COVID-19 and those deemed as healthy.

Preliminary results on the data above with a simple 3-layer convolutional neural network with binary classification showed the ability to identify COVID-19 with precision of 75% and recall of 80%, resulting in an overall accuracy of 78%. The only pre-processing methods used for these results was to normalize the images to 300 pixels in width by 300 pixels in height. Refer to Figure 1 for the charted results of validation accuracy against training accuracy.

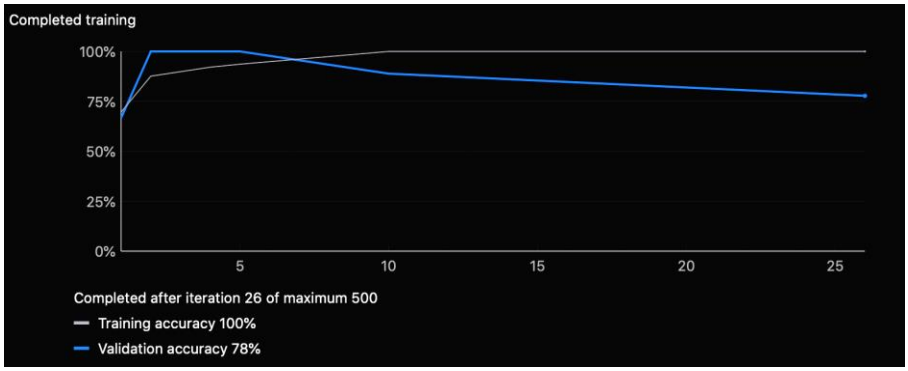

Figure 1 - Version 1 of preliminary results distinguishing COVID-19 from healthy X-ray images. It can be seen how training and validation accuracies continue to diverge, hinting at a need for change in network architecture.

Using a nine-layer network, further accuracy was achieved. Still, with no pre-processing performed on the data other than size normalizing, the model

achieved precision of 89% in F1 score against 80% recall. Further insight is offered below in Figure 2**.** Although the graphs both below and above show training accuracies of 100%, we have reason to hold this suspect. Furthermore, the algorithm claims to achieve 100% accuracy in learning the training data, but those numbers were not achieved in validation. In general, presentable numbers on algorithm precision involve those associated with the validation.

Following these findings, the analysis team began incorporating more data over a deeper network. The next two networks utilized 16 layers and 19 layers presenting precisions of 93.4% and 95.0% respectively with recalls of 89.2% and 95.1% respectively. These results are summarized in Figure and Figure 4**.**

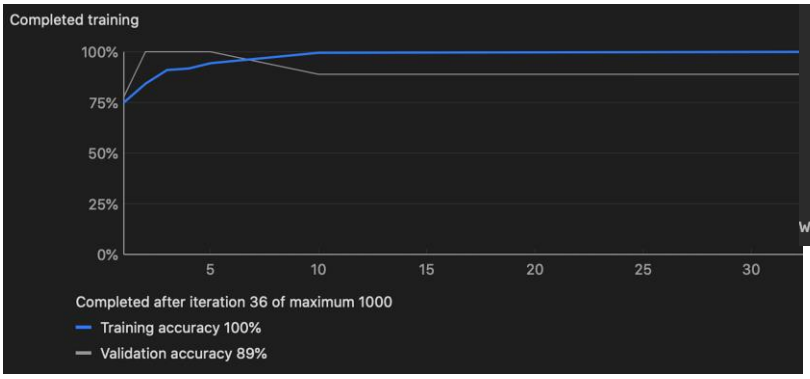

	precision	recall	f1-score
covid	1.00	0.80	0.89
normal	0.83	1.00	0.91
micro avg	0.90	0.90	0.90
macro avg	0.92	0.90	0.90
weighted avg	0.92	0.90	0.90

Figure 2 - Version 2 of the model utilizing a 9-layer network.

	precision	recall	f1-score
Covid19	0.93	0.89	0.91
Normal	0.90	0.94	0.92
micro avg	0.91	0.91	0.91
macro avg	0.91	0.91	0.91
weighted avg	0.91	0.91	0.91

Figure 3 - Results from a 16-layer network

	precision	recall	f1-score
Covid19	0.95	0.95	0.95
Normal	0.95	0.96	0.95
micro avg	0.95	0.95	0.95
macro avg	0.95	0.95	0.95
weighted avg	0.95	0.95	0.95

Figure 4 - Results from a 19-layer network

Upon these results, the team gained confidence in the direction of the architecture and began incorporating data from all 7 classes of the dataset. The final architecture utilized was a 21-layer VGG-style network and the team utilized this architecture for the final model in distinguishing COVID-19 from the other 6 conditions for the duration of the study. Results showed precision of 95.8% and recall of 98.1% with an F1 score of 97.2%.

Hierarchical and Transfer Learning

Nearly all deep learning algorithms emulate a fairly simple recipe by combining a specification of a dataset, a cost function, an optimization protocol, and a model. Horus-1 remained true to this process in its overall design. A multi-step process was used to teach Horus-1 how to not only identify but differentiate COVID-19 radiographic images from other respiratory illnesses, but also normal healthy non-infected lung scans. The three-step process involved deep learning, hierarchical learning, and transfer learning procedures. A significant image pre-processing procedure was implemented later on to optimize this process for real-world application.

A specific neural network architecture was constructed to significantly reduce the potential of the patient receiving a false-negative result on their X-ray image. This architecture resembles a hierarchy and essentially creates a “network of networks” that requires an input image to be input into several subsequent convolutional networks before it can be eventually ruled as “positive” or “negative” for COVID-19. The architecture is designed to slightly favor recall over precision due to the severity of a false-negative occurrence in this particular application. Each network in the hierarchy follows a VGG-style architecture for a total of 7 different convolutional neural networks. **Figure 5** illustrates this architecture. Each of these networks contain at least 9 layers with the largest network containing 21 layers. Each of the 7 networks

was trained and evaluated on a different permutation of the 7 conditions and then organized into a hierarchical structure.

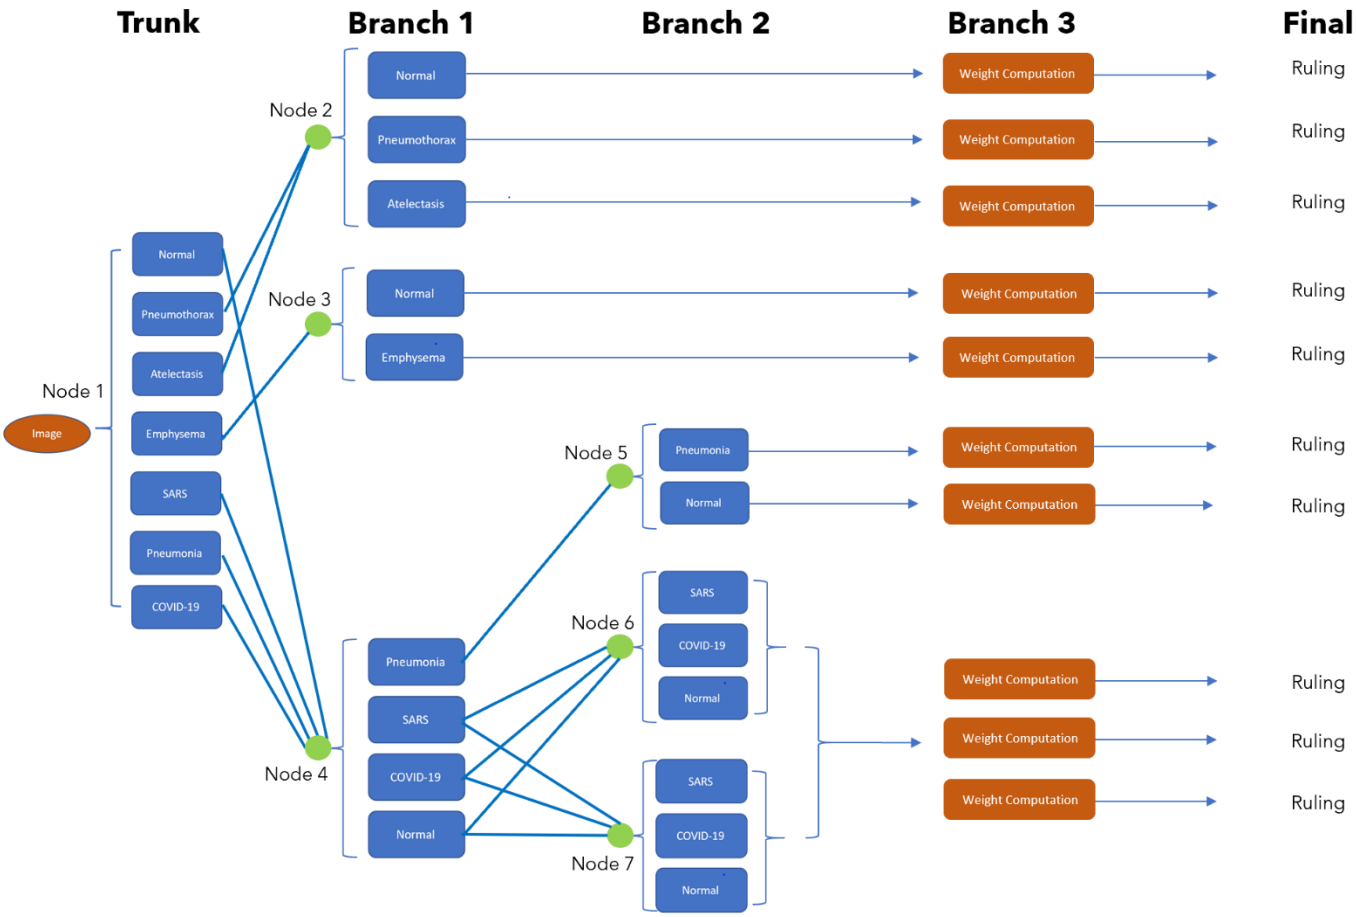

Figure 5 - Horus 1 Hierarchical Network. Each node represents a separate CNN.

As an image is first classified in Horus-1, it enters a CNN containing output labels of all 7 different conditions; this first CNN is denoted as the *Trunk* in Figure . Contingent on which classification is determined from this trunk network, it is then fed into a subsequent branch network in Node 2. Similarly, based on the classification from the first branch network, it enters into either another CNN in a subsequent branch or moves to a final weight comparison that used to determine the final ruling. This final weight comparison is denoted in *Branch 3* in Figure . At this point, scores taken from each of the CNNs in the experienced nodes are evaluated for a final classification result.

Each node of the network contains a CNN that outputs a confidence in classification that is fed into subsequent nodes for use in the final prediction. In order for an image to gain a true positive result of COVID-19, the image must go through at least 4 nodes of convolutional neural networks and a post-statistical weighted analysis. Only after this process is the final result presented. This helps maximize recall and minimize false negative results due to the fact that there is a consistent “check” against multiple networks. In the event that X-ray image is classified as “Normal” in either nodes 1, 2, or 3 of Branch 1, it is subsequently fed into Node 4 of Branch 1 of the hierarchical network for further false-negative check that the image indeed does not represent either SARS or COVID-19. Nodes 6 and 7 of Branch 2 contain two networks that classify the same three categories of the image simultaneously and a statistical average of the findings are evaluated from the results of both networks. For example, if the first network in Node 6 outputs a 65% confidence in the image representing COVID-19 and the second network of Node 7 outputs a 98% confidence in the image representing a healthy individual, the image is not automatically deemed as

"healthy." The confidence numbers and weights from the networks of all previous nodes are compared against these new numbers to contribute to the final ruling; if a ruling still receives an overall low-confidence, it is fed back into the network for reclassification until a ruling is achieved above a threshold of 96%. As this project matures, more nodes may be added insomuch that statistical error is not seen to propagate from node to node so heavily that the hierarchical architecture is invalidated by design.

Consider the mathematical model in Equation 1 of an individual neuron of a simple convolutional neural network:

$$f(\mathbf{x}) = b + \sum_{j=1}^{N_H} v_j h(\mathbf{x}; \mathbf{u}_j)$$

Equation 1

The network takes an input ***x*** and contains one hidden layer with N_H units, where v_js represent the output weights and *h(**x**;**u**)* is the transfer function dependent on those weights. A linear combination of the outputs of the hidden units with a bias *b* is used to obtain *f(x)*[11]. In a similar manner to how the weights, ***u***, and biases, ***b***, are used to influence the final output of the neuron in a standard CNN, weights and biases are calculated and analyzed at each network node to influence the final outcome of the hierarchical network after *Branch 3.*

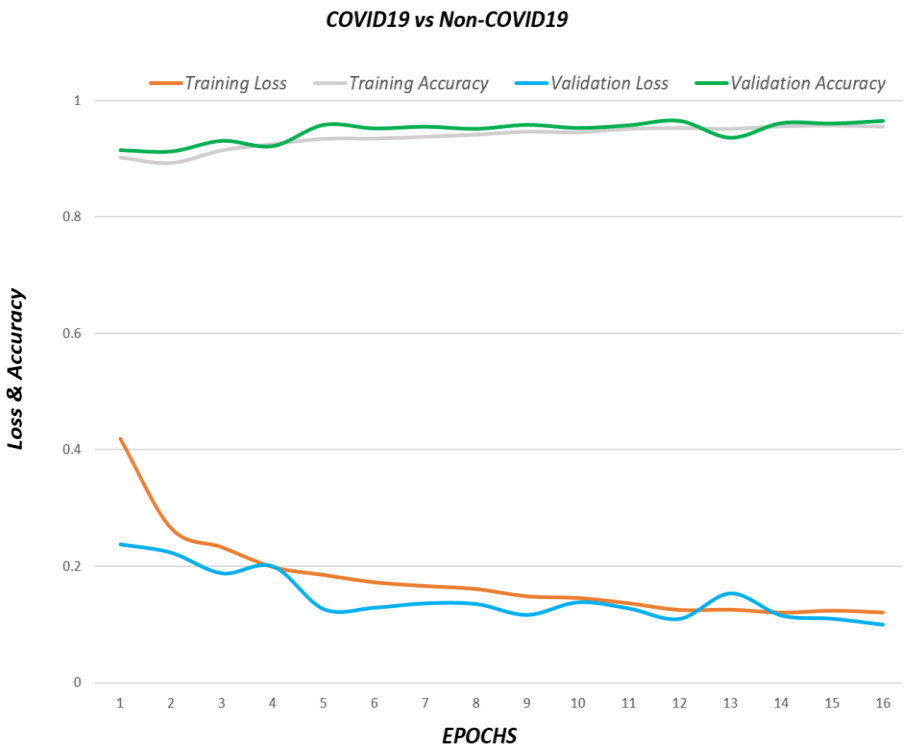

Figure 6 - Training loss plotted against training accuracy for the fourth and final CNN experienced by a true-positive COVID-19 image.

This architecture was evaluated on the mentioned dataset on a desktop computer with an external graphics processing assembly to achieve results on cleanly cropped images from the mentioned datasets. Inputs into the network were first normalized since image dimensions of the data were not all identical. To do this, the radiographic images were all rescaled to a size of 300 pixels in width by 300 pixels in height in 3 dimensions along the RGB channel. As an example, the results are shown in **Figure 6** for the final evaluated Node 7 network before it was compressed for the use of edge-computing. These results illustrate the convergence to a recall of 98.1%.

The network was then compressed to a size small enough to fit within an app on an everyday tablet or smartphone; in this case the application developed was for the iOS operating system for an Apple iPhone or iPad device. This network was the final model developed for the "import" feature within the application.

To the knowledge of the authors, with the results above and many of similar published efforts in the same field as this paper, the machine learning models received generally clean and cropped radiology This

approach works significantly well when images are imported into the model for classification. However, if a radiology technician or patient was to use the app to capture an image of the radiology image through the on-device camera, it is inevitable that several optical fallacies will exist within their image, such as glares, background scenery, reduced lighting, and general misalignment. Due to this, a significant amount of data pre-processing was needed for this effort to de-noise the data and normalize the images input into the network as much as possible.

In part, this was accomplished through the implementation of the hierarchical design. In addition, the app made use of a live object detector that identifies a chest radiographic image within the live camera feed; the frame rate of the camera feed was 60 frames per second. The architecture, algorithm, and approach behind the object detector is outside the scope of this paper, but partially utilized the Histogram of Gradients[3] (HOG) approach for detection. This object detector automatically provides locations of the corners of a bounding box around the image, so the image can be automatically cropped to the dimensions of that box. Since our neural network receives an input image of 300 x 300 dimensionality, the bounding box crops the image to a 300 x 300 resolution. Once the radiographic image is identified and cropped, optical properties are then autonomously adjusted to realize more image resolution; these properties include exposure, saturation, vignetting, and focus. Consequently, the resulting image is rid of much of the noise that was present in the original image and it can now be fed into the neural network. Figure shows the network layout incorporating this pre-processing.

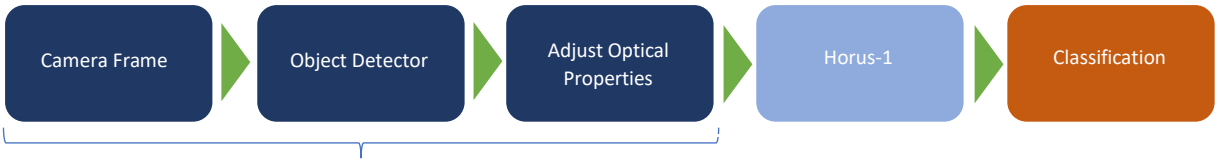

Figure 7 - This diagram explains the pre-processing procedure used.

The effects of incorporating the pre-processing module into the algorithmic flow of the operation have been astoundingly positive in field testing. However, due to the nature of field testing in a practical setting, the effects have yet to be explicitly quantified and will be evaluated in a future publication by AI Biokinetic Technologies.

Importance of Bayes Theorem in a Hierarchical Network

Bayes theorem is a foundational element to every statistical learning algorithm. Concisely, it decomposes the probability of a specific event being achieved by basing it on prior knowledge of the contributing conditions of that event. It is expressed mathematically in Equation 2[11].

$$p(\mathbf{y}_A|\mathbf{y}_B) = \frac{p(\mathbf{y}_A)p(\mathbf{y}_B|\mathbf{y}_A)}{p(\mathbf{y}_B)} \qquad \textit{Equation 2}$$

The hierarchical architecture in Horus-1 utilized Bayes Theorem heavily between nodes in order to compute the final ruling of scores and weights accumulated between those nodes. More explicitly, at the conclusion of each network node, class-conditional distributions $p(x|y)$ for $y = C_1, \ldots, C_C$ with the prior probabilities of each class were modeled. These values were then used to compute the posterior probability for each class using Equation 3[11].

$$p(y|\mathbf{x}) = \frac{p(y)p(\mathbf{x}|y)}{\sum_{c=1}^{C} p(\mathcal{C}_c)p(\mathbf{x}|\mathcal{C}_c)} \qquad \textit{Equation 3}$$

For example, referring to Figure 5 - Horus 1 Hierarchical Network. Each node represents a separate CNN.Figure , the probability of a radiographic being classified as COVID-19 at the outcome of *Branch 1*, given that it has been classified as COVID-19 through the trunk of the network can be expressed below in Equation 4.

$$P(COVID\text{-}19)_{Node\ 1\ CNN\ Prior} = 0.97$$
$$P(COVID\text{-}19)_{Node\ 4\ CNN\ Prior} = 0.98$$
$$P(COVID\text{-}19)_{Node\ 4\ |\ Node\ 1} = 0.95$$

Equation 4

$$P(COVID\text{-}19)_{Nodes\ 1\ and\ 4\ Posterior} = (0.97 \times 0.95) / 0.98$$
$$= 0.94$$

A similar process can be followed for each classification-node permutation encountered by the image as it travels through the network classification sequence. Let it be re-stated that weights are attributed to these numbers and compared at the end of *Branch 2* of the network; as transfer learning occurs, these weights may shift the resulting posterior probabilities up or down.

Implementation of Transfer Learning

Transfer learning was implemented both at the beginning of model development and at the end of development. Transfer learning is a strategy where the weights and knowledge gained from a pre-trained network is transferred into another network to solve a different but related task, involving new data, which usually are of a smaller population[6]. Horus-1 utilized transfer learning in two different scenarios for this project.

The first scenario was using the knowledge gained in previously built neural networks identifying confirmed Common Bacterial Pneumonia (CBP), Emphysema, Pulmonary Nodules (PN), Atelectasis, Pneumothorax (PTX), and healthy conditions within radiographic x-ray images of the chest. The pre-trained network was used to learn representations useful to quickly generalize overall form and generic image properties standard to radiographic images[10]. From this, the network can focus on learning biomarkers specifically indicative of COVID-19 and SARS. Through this, the time for training is substantially reduced and a network architecture that has been proven and validated in the past in a similar application can be leveraged into another one.

The second scenario utilizing transfer learning took place during field testing. In a different, yet still supervised manner of learning, users of the mobile application would import and capture images of validation data and input the correct classification when Horus-1 output the incorrect one. This allowed an important albeit minute adjustment of the network weights to further increase accuracy while also allowing us to alter the network to accommodate optical anomalies common to practical use.

Results

Based on the results of this proof of concept study, it has been demonstrated that the artificial intelligence algorithms constituting the final *Horus-1* architecture has significant effects on the automatic early detection of the novel COVID-19 virus in the palm of the user's hand.

Limitations found during the laboratory study of importing radiology images most significantly included the lack of data. A very limited amount of COVID-19 radiographic imagery data is available to the public. However, incorporating transfer learning with data augmentation techniques and pre-processing allowed the engineering team to achieve a level of accuracy in the model that showed to converge towards numbers

typically represented by larger datasets. Many attempts to challenge Horus-1 were made by introducing a variety of different infected radiographic images of the lungs, but Horus-1 was able to quickly discern the other diseases from COVID-19 patients on average 96% of the time. It is anticipated that the sensitivity and specificity of Horus-1's ability will continue to improve as more radiographic images are fed into its neural network to analyze. Through a novel network architecture incorporated with many foundational machine learning techniques, Horus-1 was able to converge to high accuracies quickly and identify biomarkers in the lungs specifically indicative of the COVID-19 virus.

For captured images with an edge device, limitations experienced were on the discontinuity with model accuracies on laboratory images versus captured images with an edge device in the field. Accuracy would be validated at some high percentage in the lab with nicely cropped radiology images, but then accuracy on the same model would not be achieved in a practical setting where the network classified images captured by the user on edge devices (in this study, those devices were smart phones and tablets).

Besides, the advantage of automatic and fast biomarker detection of an infected patient with COVID-19, the technician obtaining the radiographic image of the patients' lungs can be far removed from the potentially infected patient thereby reducing their risk of becoming infected.

This study appears to show the possibility of a low-cost, rapid, and automatic biomarker identification of the COVID-19 disease all on a handheld device. Horus-1 may empower clinicians and security personnel to more efficiently screen for COVID-19 infected patients in airports, at both the federal and state borders, and anywhere a mobile X-ray machine can be placed. The authors hope that the results presented in this publication inspire institutions to realize the power of remote diagnostic techniques achievable through mobile technology.

Next Steps

The authors are aggressively training the Horus-1 in both laboratory and practical settings to achieve higher percentages of recall; the authors plan to release results contingent on further findings in a future paper. Along with the classification of radiographic images, it would be advantageous to the end-user and/or clinician to have multiple angles of identifying symptoms in order to more accurately screen symptomatic individuals for COVID-19. To accomplish this, more analysis tools are being developed by AI Biokinetic Technologies to integrate directly into the mobile application hosting Horus-1. These tools include an audio analysis feature that identifies respiratory condition of the individual through various audio recordings, a contactless infrared temperature sensor, and diagnostic questionnaires. The covariance model of these data mediums in relation to each other have the potential to swiftly identify the probability an individual is subject to or is currently experiencing symptoms of COVID-19.

Efficacy of these tools are currently being analyzed and their feasibility in a clinical setting validated. The potential of the Horus-1 artificial intelligence platform along with the use of these tools in screening for conditions other than COVID-19 are also being heavily evaluated.

Since artificial intelligence also resembles a "black box" in nature to its algorithmic design, it would be highly relevant to incorporate explain ability into the Horus-1 platform to understand why certain classifications were made and to easier rectify false readings. This is addressed, in part, through the use of a hierarchical architecture, but the authors are contributing further efforts into creating thorough explain ability throughout the entire platform. Many explainable protocols have been proposed[12] and their efficacy evaluated, but a truly novel approach will be needed in order to address sensitive medical applications.

Discussion

Based on this proof of concept study's results, it is demonstrated that deep learning within Horus-1 may have significant and positive effects on the quick and automatic detection of biomarkers or signatures of COVID-19 using radiographic images within approximately four seconds after upload.

It is hypothesized that as more images are fed into Horus-1's neural network the sensitivity and specificity will continue to improve. Additionally, by adding Covariate data such as a patient questionnaire of symptom history, using the infrared scanner of the hand-held device to measure patient presentation and temperature, and collecting vocal samples of the patient may also make Horus-1 more sensitive and specific in not only identifying biomarkers but diagnosing patients with COVID-19 infections.

References

1. Meriam Webster Dictionary. https://www.merriam-webster.com/dictionary/COVID-19. Accessed 10 Apr 2020.
2. Smith G (2018) The AI Delusion. Oxford University Press https://www.ubs.com/microsites/artificial-intelligence/en/new-dawn.html Accessed 20 Mar 2019
3. *Artificial Intelligence: A Modern Approach. Third Edition*. Russel, Stuart J.; Norvig, Peter. 2015, Pearson India Education Services Pvt. Ltd. p 961-962.
4. Xu Z, Shi L, Wang Y et al (2020) Pathological findings of COVID- 19 associated with acute respiratory distress syndrome. Lancet Respir Med. https://doi.org/10.1016/S2213-2600(20)30076-X
5. Tech Times News. Coronavirus Has THREE Distinct Strains, According to Study; US Suffering From Original Variation. Extracted 11 Mar 2020. https://www.techtimes.com/articles/248721/20200410/coronavirus-has-three-distinct-strains-according-to-study-us-suffering-from-original-variation.htm
6. Weiss K, Khoshgoftaar TM, Wang D (2016) A survey of transfer learning. J Big data 3:9
7. *COVID-19 Image Data Collection.* Cohen, Joseph Paul; Morrison, Paul; Dao, Lan. https://github.com/ieee8023/covid-chestxray-dataset. 12 April 2020.
8. *NIH Chest X-ray Dataset.* National Institute of Health. https://www.kaggle.com/nih-chest-xrays/data/version/1. 8 October 2019.
9. *COVID-CT Dataset: A CT Scan Dataset about COVID-19.* Zhaou, Jinyu and Zhang; He, Xuehai; Xie, Pengtao. https://github.com/UCSD-AI4H/COVID-CT. 12 April 2020.
10. *Deep Learning.* Goodfellow, Ian; Bengio, Yoshua; Courville, Aaron. 2016, Massachusetts Institute of Technology. p 147 -149, 525 – 527.
11. *Gaussian Processes for Machine Learning.* Rasmussen, Carl Edward; Williams, Christopher K. I. 2006, Massachusetts Institute of Technology. p 34 - 35, 90, 199 - 200.
12. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning.* Samek, Wojciech; Montavon, Gregoire; Vedaldi, Andrea; Hansen, Lars Kai; Muller, Klaus-Robert. 2019, Springer Nature Switzerland AG.