Solution Brief

DarwinAI
COVIDNet-CT Neural Network Architecture

DarwinAI's Deep Learning AI Screening Tool Helps Detect COVID-19

DarwinAI developed the COVID-Net convolutional neural network (CNN) architecture to help clinicians detect COVID-19 in patients, with additional optimizations made using the Intel® Distribution of OpenVINO Toolkit

Author

Lauri Minas
Edge AI Marketing
Intel IOTG

"DarwinAI has been working with Intel for years to accelerate Artificial Intelligence on Intel hardware. Intel has been additionally supportive of our work on Explainability, which ensures that enterprises can adopt transparent and trustworthy AI at scale."

- Sheldon Fernandez
CEO of DarwinAI

Table of Contents

Introduction

The COVID-19 pandemic has thrown the world into an unfamiliar and uncertain place. The pandemic continues to have a tremendous impact on patients and healthcare systems around the world. In the fight against this novel disease, there is a pressing need for rapid and effective screening tools to identify patients infected with COVID-19.

As the COVID-19 pandemic took hold, the medical community reported that one of the largest bottlenecks in triage and diagnosis was the scarcity and processing time of the standard reverse transcription polymerase chain reaction (RT-PCR) viral test. In response, DarwinAI collaborated with researchers at the University of Waterloo's Vision and Imaging Processing Lab to develop COVID-Net CT as a complementary tool to assist clinicians in rapidly screening for COVID-19 and assessing disease progression and severity.

Challenges

Rapidly distinguishing COVID-19 infections from abnormalities caused by other lung conditions

In the fight against COVID-19, there is a pressing need for fast and effective screening tools to identify patients infected with COVID-19 to ensure timely isolation and treatment. Currently, reverse transcription polymerase chain reaction (RT-PCR) testing is the primary means of screening for COVID-19, as it can detect SARS-CoV-2 ribonucleic acid (RNA) in sputum samples collected from the upper respiratory tract. But RT-PCR testing is a time-consuming process which is in high demand, leading to possible delays in obtaining test results.

Early studies on CT-based screening have reported abnormalities in chest CT images which are characteristic of COVID-19 infection, but these abnormalities may be difficult to distinguish from abnormalities caused by other lung conditions.

For radiologists, visual analysis of CT scans is a time-consuming manual task, particularly when patient volume is high. Besides accelerating and improving processes, the application of AI methods has the potential to improve outcomes and reduce costs in healthcare. This has never been so relevant as during the COVID-19 pandemic.

The Solution

Human-machine collaborative design with COVIDx-CT dataset

To start building the COVID-Net CT, a dataset was derived from 104,009 images across 1,489 patient cases.[1] The raw data was cleaned and prepared in a format suitable for benchmarking, as well as providing bounding box annotations for the body regions within the chest X-ray or CT images.

> COVID-Net CT design employs a diverse collection of architectural traits that result in a high-performance model purpose-built for making accurate COVID-19 detections based upon chest X-ray images or CT scans.

After building the dataset, now named COVIDx-CT, the team at DarwinAI then used a human-machine collaborative design strategy in which they combined human-driven network design prototyping with machine-driven design exploration over four steps:

Step 1: Network Design Prototyping. After data collection to build the machine learning models, the first stage of the human-machine collaborative design strategy is principled network design prototyping. The prototype provides the framework of the model while leaving final microarchitecture and micro-architecture decisions to the machine-driven aspect of the process.

To help clinicians' better triage as well as decide on the treatment strategies, the prototype was designed to make one of three predictions:

1. No infection (normal)
2. Non-COVID-19 infection (e.g., non-COVID19 viral, bacterial, etc.)
3. COVID-19 viral infection

From this point, the initial prototype needed to be converted into a deep neural network model.

Step 2: Machine-Driven Design Exploration. The second stage of the design was machine-driven design exploration using DarwinAI's GenSynth platform that uses AI to build AI models. Employing traditional forms of machine learning, the platform observes a neural network and then uses those observations to build new, optimized versions of that network. GenSynth can reduce the complexity of designing high-performance deep learning solutions and generates highly optimized models suitable for edge computing.

GenSynth is complementary to the Intel® Distribution of OpenVINO™ toolkit for model optimization, an open-source platform that speeds deployment of deep learning applications on Intel® architecture.

For COVID-Net CT, the operational requirements included greater than 80 percent sensitivity and positive predictive value (i.e., probability that a positive prediction is in fact correct). These parameters were chosen to enable the platform to strike an appropriate balance between accuracy and speed. In addition, a goal was set to design a neural network model that could run on different platforms — be it in the cloud or on an edge device (perhaps even the actual imaging device itself).

The team at DarwinAI recognized that some hospitals might not have access to a high-performance computing datacenter to run their model. For inference/runtime, DarwinAI tested COVID-Net CT on both Intel® Core™ i5 processors as well as Intel® Core™ i7 processors. COVID-Net CT performed very well on both processors demonstrating that COVID-Net CT could be run on edge computing devices in hospitals.

Step 3: Validation via Explainability. When building a precise and robust neural network model, it is important to recognize if it is producing the right results for the right reasons. The opaque nature of deep learning is increasingly being scrutinized as AI is becoming more widely used. In healthcare, this level of opaqueness makes it difficult not only to design networks, but also to gain widespread adoption with clinicians.

The GenSynth platform automatically groups different error scenarios to provide a quick high-level picture of how the network is performing. GenSynth provided detailed information into the different error scenarios to pinpoint specific biases and gaps, as well as understand the critical factors behind model decisions as illustrated in the images below.

Ensuring that a neural network is making the right decisions for the right reasons is an essential part of designing robust models for real world applications. DarwinAI's explainable AI (XAI) technology provides unparalleled insights, highlighting the critical factors in the decision-making process so designers can identify and remove false cues from the model.

Given the data-driven nature of deep learning, 'right decision for the wrong reason' scenarios are not uncommon and can be extremely difficult to track and identify without XAI-driven auditing. In the healthcare industry, the value of explainability in improving the reliability of deep neural networks for clinical applications cannot be understated.

Step 4: Explaining and understanding the COVID-Net CT architecture. By fusing human domain knowledge with the capabilities of GenSynth, the DarwinAI team produced an effective model in seven days. Compared to the popular ResNet-50, COVID-Net CT had less than half the computational complexity while having 8% higher sensitivity for the task of COVID-19 detection[1].

DarwinAI expanded COVID-Net CT with a version that has a much larger dataset of chest CT scans. This was identified as COVIDNet CT-2. This new model exhibits an efficient micro-architecture design composed largely of 1x1 convolutional layers and depth-wise convolution layers. The heavy use of a projection-expansion-projection (PEPX) design pattern facilitates formidable performance efficiencies, while still maintaining strong COVID-19 sensitivity and positive predictive value.

> *"With regards to open source initiatives like COVID-Net CT, there's a significant cost-benefit for practitioners and hospitals to leverage existing hardware and run the system on CPUs. To this end, clinicians and academics have realized substantial efficiency gains by deploying the system through OpenVINO, and we've been thrilled to partner with Intel in making COVID-Net CT available to front-line and hospital workers using CPUs. Our collaboration has been one of the highlights during these challenging times."*
>
> \- Sheldon Fernandez, CEO of DarwinAI

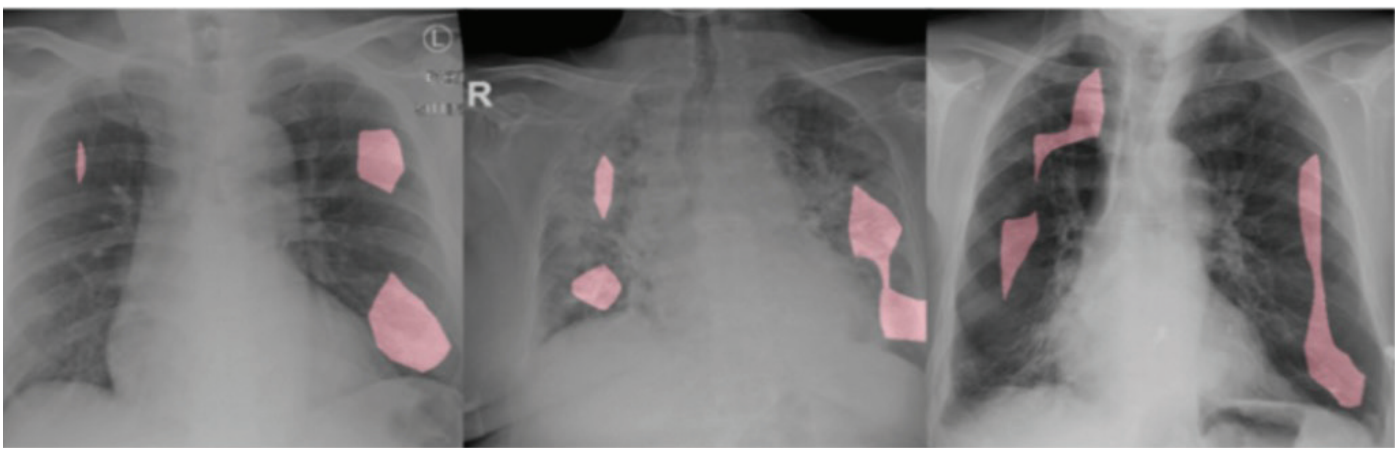

Figure 1. Example CXR images of COVID-19 cases from different patients and their critical factors in red, as identified by GenSynth

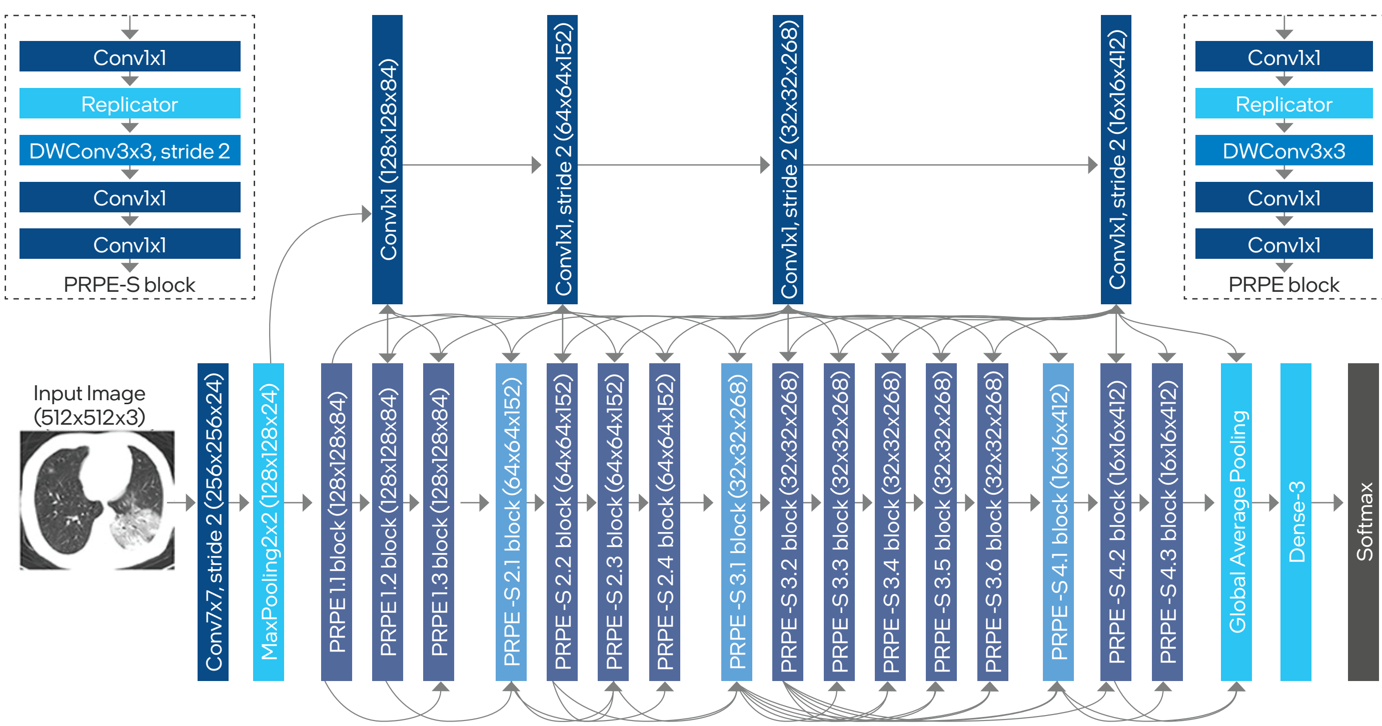

Figure 2. COVIDNet CT-2 architecture design via machine-driven design exploration includes high architectural diversity, selective long-range connectivity and lightweight design patterns.

As shown in Figure 1, the use of lightweight design patterns, such as PRPE and PRPE-S enables COVIDNet CT-2 to achieve high computational efficiency while maintaining high representational capacity. While these design patterns may be difficult and time-consuming to design manually, machine-driven design allows for these fine-grained design patterns to be rapidly and automatically discovered. Finally, selective long-range connectivity can be observed, which enables greater representational capabilities in a more efficient manner than densely connected deep neural networks.

Results

High accuracy, low compute requirements

DarwinAI's COVIDNet CT-2 neural network model architecture achieved a test accuracy of 98.1%, while having relatively low architectural and computational complexity.[1] This highlights one of the benefits of leveraging machine-driven design exploration for identifying the optimal macroarchitecture and microarchitecture designs for building a deep neural network architecture. In the case of COVIDNet CT-2, the result is a highly accurate yet highly efficient deep neural network architecture that is suitable for scenarios where computational resources are a limiting factor. In clinical scenarios, such architectures could be suitable for use in embedded devices such as the imaging systems.

The performance of the first COVIDNet CT dataset was compared with existing deep neural network architectures for the task of COVID-19 detection from the initial dataset of chest CT images. More specifically, COVIDNet CT was compared with three state-of-the-art deep neural network architectures: ResNet-50, NASNet-A-Mobile, and EfficientNet-B0.

It can be observed from Table 1 that COVIDNet CT achieves the highest test accuracy and lowest architectural complexity amongst the tested deep neural network architecture.

Going Forward

DarwinAI has recently launched COVIDNet CT-2 as an open source neural network architecture for COVID-19 detection, and has made their COVIDx-CT image dataset publicly available. DarwinAI researchers say that the scarcity of COVID-19 radiography images in the public domain was a main motivation for their open-sourcing the COVID-Net CT-2 project, as access to training data is key to boosting the speed and accuracy of AI-powered CT lung scan-based diagnosis.

Conclusion

Accelerating COVID-19 identification with an AI solution optimized by Intel® technology

The global crisis brought on by COVID-19 has affected us all. Over the last year AI researchers have worked in numerous ways to develop solutions that can help radiologists interpret chest radiography images with greater speed to accurately diagnose COVID-19 in patients. Intel processors and the Intel Distribution of OpenVINO Toolkit have been key foundational tools in DarwinAI's development of their COVID-Net CT deep neural network architectures. These AI-powered predictive models work in a way that satisfies sensitivity and positive predictive value requirements, while also minimizing computational and architectural complexity to enable widespread adoption in clinical environments where computing resources may be limited.

DarwinAI's internal testing found that COVIDNet CT-2 is capable of quickly identifying COVID-19 infections with a 98.1% accuracy rate.[1] This allows radiologists the ability to diagnose greater number of patients. COVIDNet CT and COVIDNet CT-2 neural networks were trained and optimized with the Intel® Distribution of OpenVINO™ toolkit.

Table 1. Comparison of parameters, FLOPS, and accuracy (image-level) for tested network architectures on the first COVIDx-CT dataset.

Architecture	Parameters (M)	FLOPs (G)	Accuracy (%)
ResNet-50 (25)	23.55	42.72	98.7
NASNet-A-Mobile (29)	4.29	5.94	98.6
EfficientNet-BO (30)	4.05	**4.07**	98.3
COVIDNet-CT	**1.40**	4.18	**99.1**

Best results highlighted in bold.

"By building COVIDNet CT-2 from such rich and voluminous data, we've been able to achieve a new level of accuracy of COVID-19 detection, at 98.1% with a very large and diverse dataset, and overall sensitivity now exceeds 96% across a wide and diverse number of scenarios."

- Sheldon Fernandez, CEO of Darwin-AI

Learn more

The Intel® Distribution of OpenVINO™ toolkit is a free toolkit for developers that accelerates performance, deep learning, and computer vision inference from edge to cloud. It supports heterogeneous processing and asynchronous execution across multiple types of Intel processors. The toolkit offers these benefits:

- Accelerates AI workloads, including computer vision, audio, speech, language, and recommendation systems
- Supports heterogeneous execution across Intel® architecture and AI accelerators using a common API for a write once, deploy anywhere efficiency
- Speeds up time to market via a library of functions and preoptimized kernels
- Includes optimized calls for OpenCV, OpenCL™ kernels, and other industry tools and libraries

Learn more about the Intel Distribution of OpenVINO Toolkit and get started using the free Intel® DevCloud test environment.

About DarwinAI

DarwinAI, the explainable AI company, enables enterprises to build AI they can trust. DarwinAI's solutions have been leveraged in a variety of enterprise contexts, including in advanced manufacturing and industrial automation. To learn more about DarwinAI, visit their website at **https://www.darwinai.com/** or follow them on Twitter, @DarwinAI or LinkedIn: **https://www.linkedin.com/company/darwinai/**

[1] Based on DarwinAI test results.

Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more at https://www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Intel technologies may require enabled hardware, software or service activation.

Intel technology is not intended for use in the diagnosis, treatment or cure of any illness, injury or disease.

0621/EOH/MESH/PDF 347387-001US