Spring 2020 / Volume 4, No. 2

Imaging Services

ANTECH IMAGING NEWS

Machine Learning in Radiology

The Future of AI in Veterinary Radiology

There is a plethora of talk in the industry today about machine learning. It sounds futuristic. It sounds exciting. It sounds 'next generation.' The reality is that we already encounter machine learning every day. Targeted advertising, 'smart' search engines, suggestions from our favorite streaming services, stock trading, and medical diagnostic tools all use machine learning in some form or another. But what exactly is machine learning vs. artificial intelligence?

Artificial intelligence (AI) is the high-level term used to refer to all areas of data science where a computer is used to simulate human decision-making capabilities. Machine learning (ML), is a subset of artificial intelligence, it uses algorithms to identify patterns in large amounts of data, without being explicitly given the rules to find the patterns. Below is another layer of machine learning, called deep learning, and a lot of the more advanced machine learning you see around you, self-driving cars, and advanced computer vision using deep learning.

In machine learning, we have to figure out how to teach the machine, and there are a few ways we can do this: supervised, semi-supervised, unsupervised, and reinforcement learning. In supervised learning, we may show the system lots of pictures we have labeled cat or dog to teach it what cats or dogs look like; eventually, the system encodes the rules it needs to use to identify cats and dogs from images. Unsupervised learning, on the other hand, is where we don't label the data. In this case, you'd give the system lots of images of cats and dogs and let it learn the differences itself ... it wouldn't know they were called cats and dogs, but it would be able to distinguish them.

Semi-supervised learning is a combination of these two things. We combine a small amount of labeled

Continued on next page

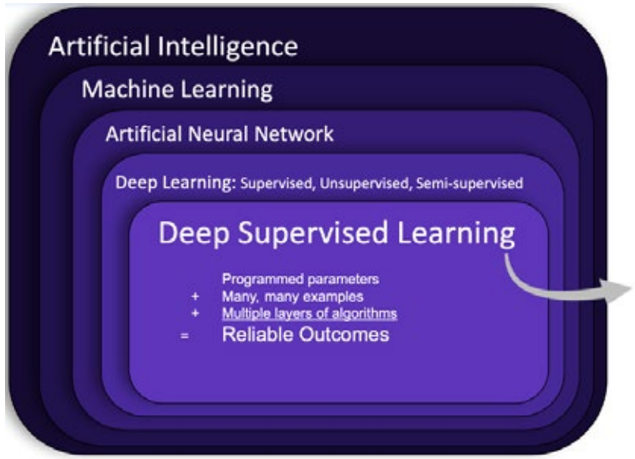

Figure 1. Diagram of artificial intelligence subsets.

data with a lot of unlabeled data. This can be extremely useful where you only have a limited amount of labeled data available to use for teaching.

Most machine learning relevant to AIS is supervised learning. Add in a deep artificial neural network plus algorithms, and now we have deep supervised learning (Figure 1.) Oh, no more vocabulary! What the heck is a deep artificial neural network? In deep learning, the deep part is derived from the fact that there are multiple layers of neural networks. These can be interconnected in a variety of different ways, depending on the type of problem the architecture is being set up to solve. But what is an artificial neural network? An artificial neural network is a program designed to mimic the decision-making process that naturally happens in humans—the natural neural network with which we already have some familiarity.

When supervised and unsupervised learning are joined, semi-supervised learning is formed. Though not discussed for the purposes of this article, unsupervised and semi-supervised learning are used in more exploratory data analysis.

The processing of multiple signals related to the fight or flight response in humans and animals is a good example of the natural neural network at work. When a human or animal is threatened, perception centers in the brain receive signals from multiple sources—eyes, ears, nose, sensory, and proprioception. The perception centers process the signals through information learned from parental teaching and past experiences. Then, a decision is made, and an appropriate response follows—fight or flight.

In the case of radiology and an artificial neural network, Roentgen signs of size, shape, number, location, margination, and opacity, all represent taught input. Many examples of a lesion are uploaded to the neural network. These examples represent the past experiences of a

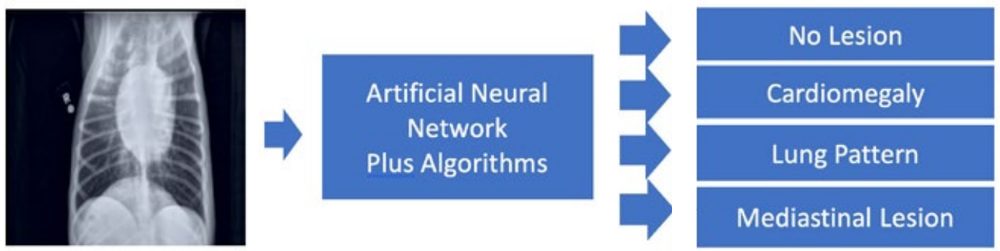

Figure 2. An artificial neural network must be equipped with algorithms to greatly enhance the accuracy of diagnosis.

natural neural network. Then examples, both positive and negative, can be uploaded, and hopefully, the machine can determine an appropriate outcome of 'lesion' or 'no lesion' (Figure 2). Hopefully.

Hopefully? If we want to remove the 'hopefully' and have the machine learn accurate outcomes, we must employ algorithms. Remember algorithms from math class? I didn't either. An algorithm is a set of rules applied to a mathematical equation. Remember the rule that says, for the equation 1+2x3, you must perform the multiplication first, then the addition? That's an algorithm. Rules applied to the equation. We apply algorithms so that the machine can learn to accurately and reliably reach the appropriate outcome.

Reliable outcomes may seem reasonably simple to achieve when broken down, as described to this point in our article. The truth of the matter is that reliable outcomes require the work of a team of expert programmers and certified radiologists. Creating the algorithms that will achieve the desired level of learning uses the best minds in programming.

The 'art' of machine learning is knowing which configuration of neural networks are best suited for the problem at hand. Too simple a structure and the model may not be able to learn the features needed to accurately make the diagnosis from a radiograph (Figure 3). Too complex, and you can create a model that's too hard to train or will over fit the data. In that case, it will only be able to expertly identify what it has already seen and will not be able to make any predictions. While deep learning is hugely powerful; it needs help getting started!

This is just one component of leveraging machine learning effectively. Having the most accurate model in the world is of no use if it takes days to process every image it looks at.

Understanding how to best leverage the data you have is a huge

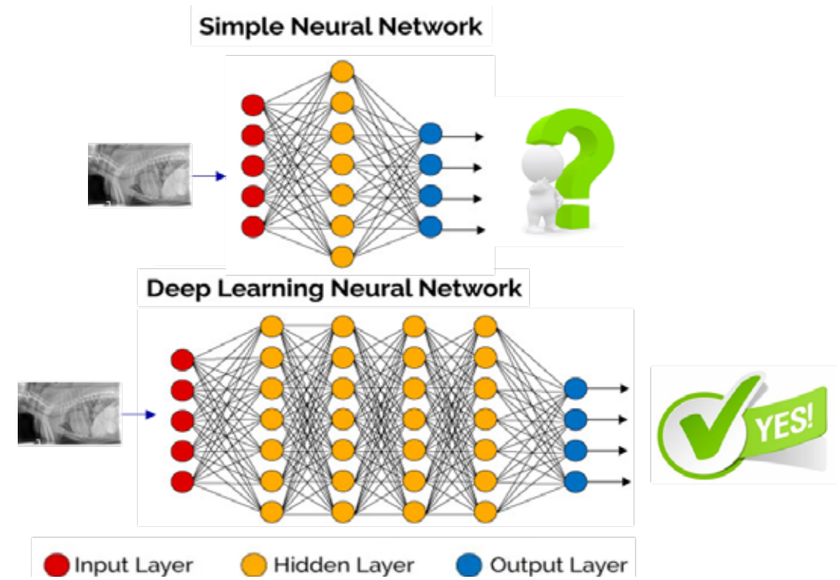

Figure 3. A depiction of a simple neural network unable to make a determination (top) and a proper deep learning neural network using multiple algorithms to reach a reliable outcome.

component of building a reliable machine learning model, and machine learning loves data! Ideally, massive amounts of data and examples are needed. And not just any data will do. The data must be verifiable. Images offered to AI as examples of radiographic findings to learn from must be verified by board-certified radiologists. Images to help AI learn to assess accurate diagnoses must be known cases with verifiable outcomes either through pathology findings, surgical exploration, or necropsy. Then, and only then, do we achieve accuracies nearing 100%.

Machine Learning at AIS partnering with The Mars Next Gen Team

Did you know that AI is already helping your Antech radiologist work more efficiently by minimizing tedious tasks and reducing the time spent between patients? This improves overall patient throughput without reducing the time or attention that our radiologists spend on each particular patient. In fact, it allows the radiologist to place more time and effort into what he or she was originally trained to do—evaluate diagnostic images.

Machine learning can also be used on the front end during image acquisition. It can be trained to help technicians by evaluating acquired images in real-time and offering suggestions for adjustments in technique, positioning, collimation, and artifact reduction. All of which contribute to higher quality images that ensure the highest quality evaluations.

One specific way AI helps is by sorting and arranging images to be read. A radiologist will typically take 3-5 seconds per case surveying the images to read, arranging the images in his or her preferred order, and righting images that are upside down, backward, or rotated. AI is used to flip, rotate, and reverse images to the correct orientation even before the radiologist accesses the case (Figure 4.) Not only that, but AI can also learn a specific radiologist's preference for how, where, and how many images are displayed as soon as the consult is accessed. Eliminating those 3-5 seconds to 'set up' the consult for evaluation means the radiologist's brainpower is better spent where it is most needed—evaluating the images.

How does AI continue to learn in order to help with differential diagnoses? (Table 1) The machine learns normal

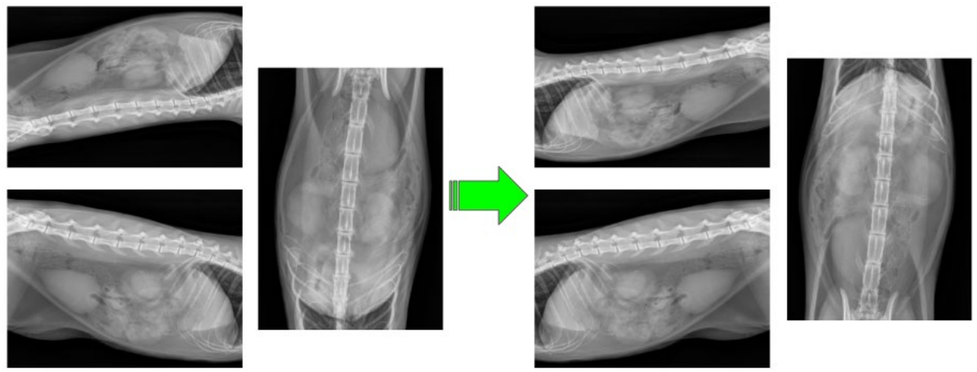

Figure 4. How AI will correct image rotation and flip as needed.

Level of accuracy (%)	Steps taken to improve accuracy
60-70%	• Initial AI programming and algorithms • Many, many examples of known outcomes input to AI
85-95%	• More known examples, radiologists' feedback through reporting, and image labeling
98+%	• Radiologists' review of AI and quality control

Table 1. Overall process for improvement of accuracy of AI.

vs. abnormal for body parts and body systems. Information such as objective measurements is readily programmed into AI from textbooks and peer-reviewed publications. At the same time, AI is learning artifacts from normal variants and verifiable lesions as well as subtle and 'hidden' lesions or lesions not readily visible to a board-certified radiologist. AI measures using known parameters. AI learns differential lists, which are also easily programmed from published textbooks and peer-reviewed publications.

AI then produces a prioritized differential diagnosis list with all imaging lesions considered and offers accurate probability based on having reviewed hundreds of thousands of cases with diplomate case reports and known case follow-up where available. Board-certified radiologists perform quality control and give direct input into a case evaluation. The referring veterinarian gets real input into the case in record time for which s/he can be confident. There is better quality of care, and more animal lives are saved.

The AIS-Mars Next Gen machine learning team is currently perfecting algorithms, and increasing diplomate verified data input to AI. We have already reinforced our deep supervised machine learning with tens of thousands of images specifically labeled by our AI team of radiology diplomats. That number of images grows daily. Accuracy is paramount. That is why we only use diplomate reviewed images and proven known cases for supervised machine learning. Currently, AIS does not market AI for case diagnosis. We will not compromise our current, near 100%, human accuracy, and AI is simply not there yet.

Beware of currently marketed, low cost, AI-assisted diagnoses that do not have board-certified radiologists performing quality control on every case. Also, beware of marketing schemes that use made-up terminology and buzzwords instead of real medical terms. In the authors' experience, such tools can determine normal from abnormal for a number of organs at a beginner veterinary student level. The accuracy of such tools is in the 50-65% range. At the time of this publication, we have not seen a veterinary product that contributes in any significant way to actual differential diagnoses.

-Diane U. Wilson, DVM, DACVR
-Mark Parkinson, Director Mars Next Gen Technologies

ANTECH IMAGING NEWS

Spring 2020 / Volume 4, No. 2

Machine Learning and Artificial Intelligence:

What to Expect from AIS in the Near Future

AIS expects to release our AI-assisted diagnostic imaging when we can guarantee machine accuracy to match board-certified radiologist accuracy – for correctly identifying true positives and true negatives and providing real input to differential diagnoses.

AI will continue to work transparently to improve the efficiency of our team of board-certified radiologists—bringing the best diagnostic imaging care possible to referring practices. Our board-certified radiologists will continue working toward 100% accuracy for AI, which, in turn, will assist diagnosis by increasing the efficiency of our boarded staff.

As we continue toward 100% accuracy for AI, additional options will be available to practices. Options will be to include the ability to have real-time educational feedback on technique and positioning for your technical staff and near-instantaneous reports directly from AI for pre-anesthetic ECGs and wellness images, and AI-assisted reports with confidence level diagnostic input (not just lesion identification). The veterinarian will always have the ability to have a human, board-certified radiologist review any study.

We do not anticipate machine learning will replace human radiologists. To us, it's an assistive technology giving people the digital armor needed for the 21st century! We do expect continued developments to allow for more efficiency in creating evaluations and higher, faster patient throughputs. This will lead to quicker returned consults, which will improve your hospital efficiency and, ultimately, contribute to higher patient care and helping animals worldwide.

Technical Support 24/7/365
1-877-727-6800
Support@AntechImagingServices.com

Contact for other questions or concerns

Lisa Ziemer, DACVR, Chief Medical Officer
Lisa.Ziemer@AntechImagingServices.com

Central U.S. Sales Representative

Ryan Reynolds, DVM
1-409-527-0408
Ryan.Reynolds@AntechImagingServices.com

Western U.S Sales Representative

Leslie Dowd, DVM
1-949-397-5338
Leslie.Dowd@AntechImagingServices.com

Eastern U.S. Sales Representative

Cory Ellis, DVM
1-803-389-7389
Cory.Ellis@AntechImagingServices.com