

Artificial intelligence in veterinary medicine

Ryan B. Appleby, DVM, DACVR^{1*}, and Parminder S. Basran, PhD, FCCPM²

¹Department of Clinical Studies, Ontario Veterinary College, University of Guelph, Guelph, ON, Canada

²College of Veterinary Medicine, Cornell University, Ithaca, NY

*Corresponding author: Dr. Appleby (rappleby@uoguelph.ca)

<https://doi.org/10.2460/javma.22.03.0093>

ABSTRACT

Artificial intelligence (AI) is a branch of computer science in which computer systems are designed to perform tasks that mimic human intelligence. Today, AI is reshaping day-to-day life and has numerous emerging medical applications poised to profoundly reshape the practice of veterinary medicine. In this Currents in One Health, we discuss the essential elements of AI for veterinary practitioners with the aim to help them make informed decisions in applying AI technologies into their practices. Veterinarians will play an integral role in ensuring the appropriate uses and good curation of data. The expertise of veterinary professionals will be vital to ensuring good data and, subsequently, AI that meets the needs of the profession. Readers interested in an in-depth description of AI and veterinary medicine are invited to explore a complementary manuscript of this Currents in One Health available in the May 2022 issue of the *American Journal of Veterinary Research*.

Artificial intelligence (AI) is an important technological advancement changing the shape of our lives. Employed by many of the largest companies in the world, AI technologies power our favorite mobile applications; offer movie, TV, and music suggestions; and predict the next word in our text messages. The opportunities to leverage the power of AI to improve our quality of life in our day-to-day or professional activities seem endless, notwithstanding ethical challenges. Nowhere is this opportunity more apparent than in medical practice. Over the last 2 decades, publications on AI in medicine have increased exponentially.¹ In particular, there have been profound advances in AI and the field of diagnostic imaging where technologies have been developed to aid in diagnosis and support radiologists in both research and commercial settings. A similar shift is occurring in veterinary medicine: we are on the cusp of a large and unpredictable shift in available technology that has the potential to reshape how veterinary medicine is practiced.

Practically speaking, it is not necessary for the average veterinary practitioner to have a working knowledge of computer programming to effectively use and implement AI. However, the implementation of adopting a technology can be considered analogous to implementing a new diagnostic test in practice. For example, the average veterinarian cannot design and develop an ELISA for benchtop testing. However, veterinarians do have enough underlying knowledge of what an antigen or antibody is, as well as how ELISA testing works to feel confident employing so-called SNAP tests in practice. Furthermore, veterinarians are aware of the need to have a quality management system in place to ensure the tests work as expected. Similarly, while details of the computer algorithm designed for diagnosis or detection of diseases will initially be foreign to most

veterinarians, a baseline level of knowledge is needed to understand the power and pitfalls of AI.

Artificial intelligence is still in its nascent stages but will likely have a profound impact on our profession in the years to come. Therefore, it is vital that all veterinarians understand both the promise and limitations of AI. The purpose of this manuscript is to provide the definitions and concepts behind AI, describe its power and pitfalls, and provide some guidance to veterinarians implementing AI in their practice.

Background

Artificial intelligence is the branch of computer science devoted to creating systems to perform tasks that would normally require human intelligence.¹ It is a broad umbrella term that encompasses a variety of subfields and techniques. While most current applications of AI have been developed in the last decade, the concepts and ideas have been around for at least 70 years. In the late 40s and early 50s, the first concepts of AI were introduced to the scientific community.² Perhaps most famously, British computer scientist Alan Turing was one of the first to introduce the concept of computers performing intelligent tasks in 1950.² The term *artificial intelligence* was coined by John McCarthy in 1955.^{2,3}

Artificial intelligence was heavily researched in the second half of the 20th century but experienced little advancement in its application and scope due to limitations in available computing power. However, with major advances in computer processing power in the past decade and the digitization and availability of large amounts of data, AI has taken off.^{1,2,4} In medicine, such data includes but is not limited to medical imaging such as radiographs, CT, and MRI; photomicrographs ob-

tained by cytology and histology; and data from medical records including free text and bloodwork results. A major application of AI in medicine is to glean insights from these massive data sets with the aid of computer algorithms to help make or improve diagnoses and improve therapy and patient outcomes. As we will see, the way in which this information can be analyzed and applied varies and depends on the desired outcome of the AI.

Types of AI

Some people picture AI as computers that talk to us or, in the providence of most science fiction, robots aiming to take over the world. However, these types of AI are not what is meant when we refer to AI for medical practice. While there are many ways to classify AI, one classification scheme relates to the scope and ability of the AI.⁵ Artificial intelligence that has humanlike intelligence is known as artificial general intelligence or strong AI. Artificial intelligence with greater than human intelligence is known as artificial super intelligence. These types of AI generate a lot of fear on the basis of their portrayal in science fiction. However, they are entirely the providence of fiction. Despite the increasing complexity and ability of computers, no known computer systems are near either general or super intelligence. In fact, some researchers argue this type of AI may never exist.⁵ In reality, AI systems are often designed for a very specific function, such as walking, talking, deciphering, and responding to verbal commands or, in the case of medicine, providing a solution to a specific clinical question.

The type of AI that exists today in medical applications is broadly considered artificial narrow intelligence. That is, these AIs are designed for a specific task and since they only do a specific task, they are considered narrow or weak. In everyday life, examples of artificial narrow intelligence are used in tasks like the spam filter on your email or predictive text for messaging and in word processors. An example in medical practice is the detection of an abnormality on a radiographic image. While these are still complex tasks that require large amounts of data (eg, thousands of radiographic images) for training the AI system, these AIs are still considered narrow or weak when viewed from the level of intelligence.

Methods and Terminology for AI

To understand how AI works and how to implement it in veterinary practice we must first understand some general concepts and terms used in AI. While there are many techniques, which are outside the scope of this article, some general terminology will be encountered by any individual reading about AI. A commonly encountered term is *machine learning (ML)*, which is a subfield of AI in which algorithms are trained to perform tasks by learning patterns from data rather than by explicit programming.⁶ It is unusual to think of computer programs as “learning,” but this is exactly what differentiates AI from rule-based computer programs.

Machine learning models can learn in 3 ways: *supervised learning*, *unsupervised learning*, and *semisupervised learning* (**Figure 1**). In supervised learning,

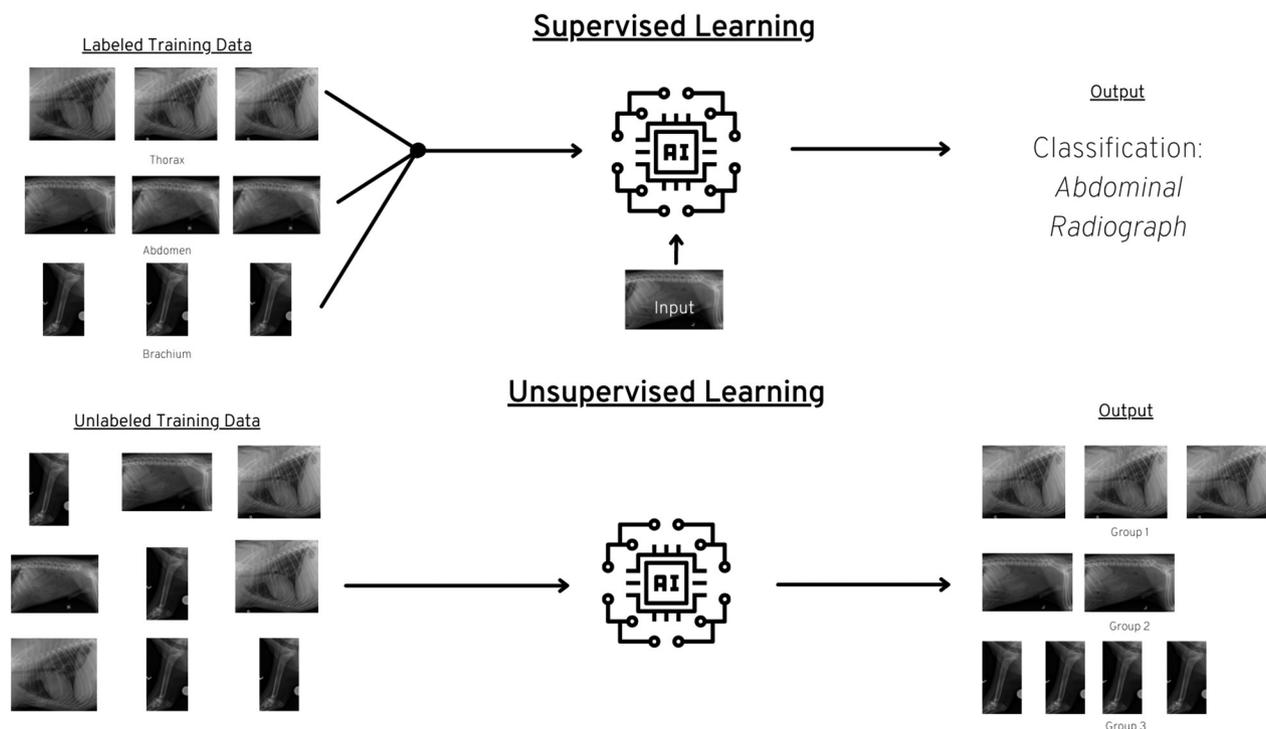


Figure 1—Examples of types of machine learning as they apply to veterinary radiology. In the case of supervised learning (top), training data in the form of radiographic images are labeled in this example into different classifications with respect to the type of study. An artificial intelligence (AI) model that encounters a novel image (input) can then classify the image on the basis of what it has learned from the training data. In the case of unsupervised learning (bottom), data without labels are grouped by the AI model into images with similar features.

labeled data sets are used to train algorithms to classify data or predict a number.⁷ Supervised learning requires the outcomes of medical data—that is, the diagnosis or classification—to be known prior to training the ML model to the task (labeled). In this type of ML, the algorithm needs 2 things: lots of data and the corresponding labels. Supervised learning is the most common form of learning for ML algorithms in medical practice as they provide the most clinically relevant data.⁷ However, unsupervised learning, where the ML algorithm generates its own set of criteria by which to classify data or predict outcomes, can be valuable, especially for large data sets in which the features to distinguish groups are unknown.⁷ In this type of ML, the algorithm has data but none of it is labeled. The goal of unsupervised learning is to help make sense of the data by means of examining it in detail and seeing if there are relationships or correlations in the data that might be clinically useful. Semisupervised learning uses a combination approach and can be valuable to developing algorithms where some of the data is missing an outcome.⁷

There are a number of ML techniques with sophisticated names like support vector machines, decision trees, naïve Bayes, logistic regression, and linear regression that rely on supervised learning. But fundamentally, these are algorithms that simply predict an outcome on the basis of some input data.

Unsupervised ML techniques include clustering and principle component analysis. Fundamentally, these algorithms find associations in data and help collapse large data sets into smaller representative ones (Figure 1). These methods are commonly considered *classical ML*.⁷ In contrast, *modern ML* involves the 2 other most common terms in AI: *artificial neural networks (ANNs)* and *deep learning*.

Artificial neural networks are computer systems composed of layers of connected nodes.³ They are named as such because they mimic the biological process of neurons, which have many input and output synapses and interconnect with other neurons with many input and output synapses. As this is an in-depth concept, it may be beneficial to consider the example of making a diagnosis of “pulmonary nodules” on a medical image both from the perspective of a human observer and an AI (**Figure 2**). In the case of the human observer, our eyes receive light signals from the computer monitor that trigger photoreceptors in the eyes. These receptors then trigger several interconnected optic nerves in the retina that send information on the shape and color of the object through the optic nerve bundles and chiasm to the brain for interpretation. Not all neurons in the retina pass signals through the exiting optic nerve to the brain, and only few neurons in the brain decode the neural messages as “medical image.” Many

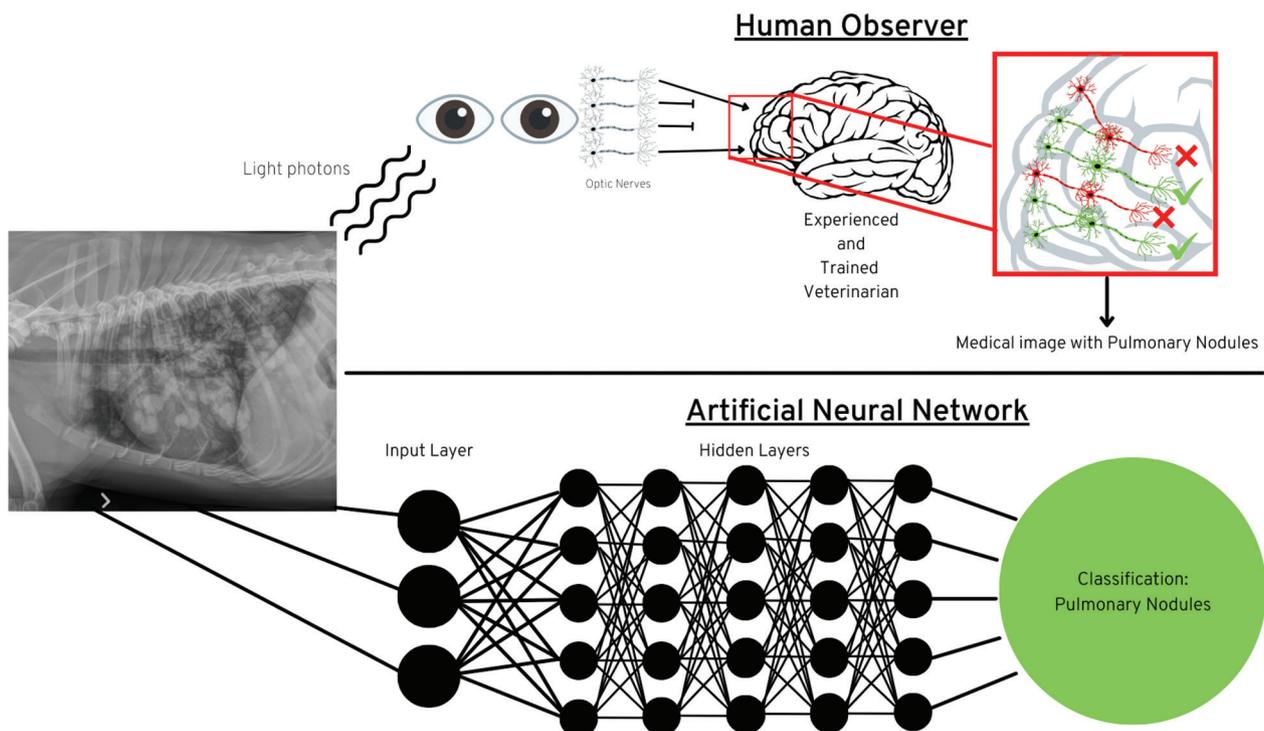


Figure 2—An illustrative example of image classification by an expert veterinary observer and an artificial neural network (ANN). In the case of a human observer, light photons trigger nerves in the retina, some of which activate and send signals to the brain. Within the brain, networks of neurons are selectively activated (green) or deactivated (red). This in turn triggers a response by a person with training and experience to classify the image as having pulmonary nodules on the basis of its similarity to other images seen previously. In an ANN, pixel data from the images enter at the level of the input layer before progressing through a series of hidden layers. Depending on the degree of activation or deactivation in the hidden layer, the ANN would classify the image, in this case correctly as a patient with pulmonary nodules.

nerves intentionally switch on and off on the basis of the size and color of the object detected. Then, the trained veterinarian interprets this image as “pulmonary nodules” since they have repeatedly seen very similar images. The training and experience of the veterinarian have developed a trained set of neural pathways that combine to generate an output of “pulmonary nodules.” The input is the shape and color of the object, and the output is the brain interpreting the input through a densely connected network of neurons switching on and off uniquely as “pulmonary nodules.”

Just as neurons receive sensory input and require a certain level and combination of interconnected activation in producing an output, an ANN consists of a network that converts input data to an output. Input data for the ANN may be the medical image. The image is processed and filtered through a series of hidden layers that help predict the output. In training an ANN, weights are applied to the hidden layers to minimize incorrect predictions.⁷ An ANN could involve just 1 hidden layer, but its value lies in the depth of layers. This is where *deep learning* comes into play.⁸ Deep learning occurs on an ANN with typically many more than 10 layers (hence, deep). This allows the algorithm to handle incredibly complex data, like medical images. Artificial intelligence for medical image analysis typically relies on a convolutional neural network, which is the most popular neural network for applications of AI in medicine.⁷

It can be challenging to understand how these terms fit together, and many are often inappropriately used interchangeably. The Russian nesting dolls analogy is one offered by IBM in the IBM Cloud Learn Hub. In this way, the concepts of AI can be thought of as Russian nesting dolls, each fitting inside one another. Machine learning is a subfield of AI. Deep learning is a subfield of ML, and neural networks make up the backbone of deep learning algorithms.

Part of the promise of AI is the opportunity to identify aspects of data that are not immediately apparent to a human observer. *Radiomics* is a field of study in which an AI algorithm can be used to extract a large amount of quantitative features from medical images.⁹ The goal of radiomics is to determine the phenotype of the imaging finding; for example, what kind of tumor is present. One can imagine the implications if a type of mass—for example, a splenic mass—could be determined from the imaging features alone or if precision-based treatments could be developed from noninvasive phenotyping of disorders.

The final important concept for veterinarians to appreciate is *natural language processing*, which is a subset of AI in which computers can decipher and attribute meaning to text and spoken words. These programs draw from ML techniques and linguistics and will be incredibly important for efficient medical practice. While medical imaging and numerical data can be directly analyzed with ML techniques, some of the most important data in medical records are text based. Natural language processing can play an important role in efficiently extracting information from medical records.

Uses of AI

The potential applications for AI in veterinary medicine are immense. Artificial intelligence can foreseeably be introduced into nearly every aspect of veterinary practice including diagnostics, companion animal care, population medicine, agriculture, research, education, and industry. If digital data exists and can be curated, AI technologies could be leveraged. In this article, our examples and considerations are limited to diagnostic medical imaging. This is due in part to the authors' backgrounds, but it is worth noting that the use of AI in clinical diagnostic imaging practice is foreseeable within the next few years, largely because much of the data (radiographs, ultrasound, CT, MRI, and nuclear medicine) and their corresponding reports are in digital form.

In veterinary diagnostic imaging, applications of AI focus on the detection, segmentation, or classification of features in the image.⁴ In the case of detection, abnormalities can be identified in images; for example, if a pulmonary nodule is present in the image. In the case of segmentation, structures can be delineated within the image; for example, defining the border of the nodule in the image. In classification, the feature or image can be assigned a category; for example, if the patient is positive or negative for metastasis.

Similar to the tasks undertaken by radiologists,⁴ detection, segmentation, and classification algorithms can be applied in other veterinary professions. These types of applications can aid in triage to ensure timely care for critical patients or timely radiology reports for images in which an AI has identified an abnormality. If AI systems can connect to digital health informatics and patient records systems, they can be used to streamline processes to reduce the workload on veterinary professionals. Or, if used in conjunction with diagnostic equipment, they can assist veterinarians in making accurate diagnoses on the basis of either the detection or classification of disease.

To date, there are a few commercially available AI systems for veterinary diagnostic imaging; however, these systems have not undergone the rigors of peer review. A majority of the peer-reviewed AI applications for veterinary imaging to date focus on proof of concept to show that AI can accurately detect abnormalities in the canine thorax.¹⁰⁻¹⁴

Considerations and Challenges

As AI becomes more available and accessible in veterinary medicine, veterinarians should consider a number of important challenges that will be encountered. We propose the following considerations for veterinarians as they consider the adoption of AI into veterinary practice.

Use cases

For AI to be most effective, it requires a directed purpose. Determining to which aspects of practice AI should be applied is a job for veterinarians. Veterinarians provide a level of expertise that determines which use cases are most valuable for the profession. Ques-

tions to consider include the purpose of the technology and how it benefits the welfare of the animal or the profession. Without veterinarians taking an active role in asking these questions when they are offered an AI-based solution, there is a risk of prioritizing profits over clinical outcomes and the well-being of both veterinary professionals and their patients. The goal of AI should be to improve veterinary practice, animal health outcomes, patient quality of life, and the lives of veterinarians. This requires well-thought-out use cases and active veterinary stakeholders.

Good data

Data is not cheap, and AI developers know this. In an era in which AI becomes a commercially viable entity, data will become like a currency. Who owns the data, how it will be used and managed, and eventually how it will be used in an AI model will be shared and should be well established at the outset of AI projects. Universities and academic institutions often have the benefit of local resources that can provide guidance on data sharing, licensing, and intellectual property; however, this may not be the case for smaller institutions and clinics. We encourage veterinary professionals to recognize their role in curating data and maintain ownership. Data must be labeled in most cases, representing a role for veterinarians that should not be underplayed.

Data is only as good as the individual labeling it and the label they apply. Labels must have an underlying ground truth. This raises the concern of what determines the truth of the label, particularly in medicine when there are often many confounding factors associated with a diagnosis or treatment. Additionally, there are differences in opinions between veterinarians or between institutions. For example, radiologists differ widely in their sensitivity to diagnosing small animal patients with a bronchial pattern. Within this range of diagnosis is a second range relating to the clinical significance of the finding. The subjective nature of interpreting medical tests creates variability in data labeling, which can lead to over- or underfitting of the AI model.

Bias is an additional consideration of concern in AI. Bias may arise from training sets that have skewed breed or geographic distributions or imposed by the method of curating the ground truth. In the case of veterinary radiology, patient positioning and radiograph quality will play a significant role in the performance and bias of algorithms. For example, AI trained on perfectly positioned and exposed radiographs may not perform well with radiographs acquired with suboptimal technique. While it is not realistic to assume data sets are perfectly labeled, careful review of data before attempting to use it in AI helps establish implicit assumptions in the model.

Not all data is useful data. It is also important to establish what data is worth properly labeling for AI use. Big data approaches in health care can be data rich but information poor (also referred to as DRIP). Indeed, some of the biggest health-care challenges are not related to what problems can be solved with AI but creating and curating good and clean data for AI.

Open data

As researchers, we believe in accessible and equitable data sharing. While there are many reasons why one should consider accessible and equitable data sharing, some are as follows. First, having access to transparent data permits an opportunity to validate data by others. Data sharing permits institutions and researchers to identify potential biases in their data collection processes before developing an AI algorithm. Second, open access allows an opportunity to validate that the data and AI system can be used on different platforms (interoperability). Data, and any AI developed from it, should not behave differently for PC, MAC, or Linux users. Third, accessible data allow different researchers and institutions to formulate different and potentially better AI solutions. Last, data sets in veterinary practice are much smaller than those found in human medicine. Combining data from different institutions allows the opportunity to increase the diversity of data sets for ML, which can eventually lead to more robust and generalizable models. It is our opinion that open data is good for veterinary medicine.

This promotes the ideal goals of AI rather than monetization. However, it must be acknowledged that there is a cost associated with data storage and sharing. This cost is likely to fall on corporations and organizations that will likely retain at least some of the rights to the data. Partnerships between academic institutions, granting agencies, and professional organizations may be necessary to facilitate open, curated, and good veterinary data for the community.

Ethics and regulation

There are many ethical questions that arise with AI in veterinary medicine. Initially, as discussed above, a question exists as to who owns the data. This is important to consider as data in the world of AI has immense value. It is also important to consider confidentiality and security as it relates to patient information.

While AI products for medical applications in humans are regulated by the FDA in the US or other similar bodies worldwide, there is no regulatory framework for AI in veterinary medicine. Regulation will be critical for the success of AI in veterinary medicine to encourage ethical, responsible, and directed use.

Implementation

While there is much promise behind AI, implementation to use AI most effectively in practice has proven to be a challenge in the human medical field. While hurdles such as FDA approval are not currently in place for veterinary medicine, the other challenges to implementation will be similar. Good AI governance practices in veterinary medicine should be developed if these technologies are routinely integrated into practice. Some critical responsibilities include establishing use cases, fiscal responsibility when AI technologies are purchased, sufficient infrastructure (software or hardware) and resource support, establishing good data management principles, thorough acceptance testing and clinical deployment including training and education, and comprehensive qual-

ity management of these systems throughout the technology's life cycle.

Client acceptance

Another aspect not to be overlooked is how clients respond to AI technologies. While somewhat outside the control of veterinarians, client acceptance of AI will be vital for its success. Veterinary professionals can promote acceptance with transparency and client education.

Conclusion

In response to a growing concern that AI would replace radiologists, in 2019 Stanford radiologist Curtis Langlotz wrote "radiologists who use AI will replace radiologists who don't."¹⁵ While there is no concern that veterinarians as a whole will be replaced by AI, given the profound potential benefits to our profession and our patients there is a likelihood that similar logic will apply and veterinarians who use AI will replace veterinarians who don't.

While there is much to consider and understand with respect to AI, there is great promise for what AI can do to improve veterinary practice. It is our hope that this article has gleaned some insight on the potential—and pitfalls—of AI in veterinary medicine and enticed you to learn more. For those interested in understanding more technical aspects of AI, we hope you enjoy the companion article in the May 2022 issue of the *American Journal of Veterinary Research*.

Acknowledgments

No external funding was used in the preparation of this manuscript. The authors declare that there were no conflicts of interest.

Figures were created using available open source image features from Canva.com.

The authors would like to acknowledge the American College of Veterinary Radiology/European College of Veterinary Diagnostic Imaging Artificial Intelligence Education and Development Committee for its role in discussions that have played a role in influencing the important features highlighted in this paper.

References

1. Tran BX, Vu GT, Ha GH, et al. Global evolution of research in artificial intelligence in health and medicine: a bibliometric study. *J Clin Med*. 2019;8(3):360. doi:10.3390/jcm8030360

2. Kaul V, Enslin S, Gross SA. History of artificial intelligence in medicine. *Gastrointest Endosc*. 2020;92(4):807–812. doi:10.1016/j.gie.2020.06.040
3. Currie G. A muggles guide to deep learning wizardry. *Radiography (Lond)*. 2022;28(1):240–248. doi:10.1016/j.radi.2021.10.004
4. Tang A, Tam R, Cadrin-Chênevert A, et al. Canadian Association of Radiologists white paper on artificial intelligence in radiology. *Can Assoc Radiol J*. 2018;69(2):120–135. doi:10.1016/j.carj.2018.02.002
5. Fjelland R. Why general artificial intelligence will not be realized. *Humanit Soc Sci Commun*. 2020;7:10. doi:10.1057/s41599-020-0494-4
6. Waljee AK, Higgins PDR. Machine learning in medicine: a primer for physicians. *Am J Gastroenterol*. 2010;105(6):1224–1226. doi:10.1038/ajg.2010.173
7. Jiang F, Jiang Y, Zhi H, et al. Artificial intelligence in healthcare: past, present and future. *Stroke Vasc Neurol*. 2017;2(4):230–243. doi:10.1136/svn-2017-000101
8. Chartrand G, Cheng PM, Vorontsov E, et al. Deep learning: a primer for radiologists. *Radiographics*. 2017;37(7):2113–2131. doi:10.1148/rg.2017170077
9. Bibault JE, Xing L, Giraud P, et al. Radiomics: a primer for the radiation oncologist. *Cancer Radiother*. 2020;24(5):403–410. doi:10.1016/j.canrad.2020.01.011
10. Banzato T, Wodzinski M, Burti S, et al. Automatic classification of canine thoracic radiographs using deep learning. *Sci Rep*. 2021;11(1):3964. doi:10.1038/s41598-021-83515-3
11. Boissady E, De La Comble A, Zhu X, Abbott J, Adrien-Maxence H. Comparison of a deep learning algorithm vs. humans for vertebral heart scale measurements in cats and dogs shows a high degree of agreement among readers. *Front Vet Sci*. 2021;8:764570. doi:10.3389/fvets.2021.764570
12. Boissady E, de La Comble A, Zhu X, Hespel AM. Artificial intelligence evaluating primary thoracic lesions has an overall lower error rate compared to veterinarians or veterinarians in conjunction with the artificial intelligence. *Vet Radiol Ultrasound*. 2020;61(6):619–627. doi:10.1111/vru.12912
13. Li S, Wang Z, Visser LC, Wisner ER, Cheng H. Pilot study: application of artificial intelligence for detecting left atrial enlargement on canine thoracic radiographs. *Vet Radiol Ultrasound*. 2020;61(6):611–618. doi:10.1111/vru.12901
14. Burti S, Longhin Osti V, Zotti A, Banzato T. Use of deep learning to detect cardiomegaly on thoracic radiographs in dogs. *Vet J*. 2020;262:105505. doi:10.1016/j.tvjl.2020.105505
15. Langlotz CP. Will artificial intelligence replace radiologists? *Radiol Artif Intell*. 2019;1(3):e190058. doi:10.1148/ryai.2019190058