Smart diagnoses: the evolving role of AI in veterinary diagnostics

Dr Richard E. Goldstein of Zoetis on how artificial intelligence is shaping the future of veterinary care

It's two o'clock in the morning. There are two dogs in my veterinary hospital's emergency waiting room and a vomiting cat in the back – anxiety is high (mine and the pet owners'). There is nobody to help but me, the on-duty veterinary surgeon.

It's stress-provoking enough just to picture the scene, I know. This is the "point-of-care" moment where we need the most effective tools at our disposal, where a timely diagnosis is imperative to reach the best possible patient outcome. On that, let me ask the question that's been on the minds of professionals in medical industries over the past year or so: (how) do we bring in additional resource to help manage workflows, while delivering on the increased client expectation on standards of care?

In the case of veterinary diagnostics, artificial intelligence (AI) has the potential to shift traditional approaches quite significantly. Analysers harnessing the power of AI typically comprise of traditional sample preparation techniques combined with newer standardised preparation methods, an automated digital scanner, and an algorithm operating within either a shallow or deep learning framework that communicates with a cloud-based system.

Machine (also referred to as shallow or superficial) learning AI is already a familiar friend to the veterinary practice; it's used every day in the form of equipment that assesses x-rays, ultrasound scans, blood analysers and other diagnostic testing devices. Systems that utilise shallow learning AI are a remarkable support in this context, but they are reliant on the manual entry of a predefined set of data represented by specific features. This leads to limitations in the system's ability to incorporate new, unseen data.

As the technology continues to rapidly evolve, deep learning AI is beginning to be integrated within the latest diagnostics devices. These deep neural networks can automatically discover important patterns in raw data, meaning these systems can identify complex representations – without the need for explicit feature engineering. The result is that it can be exposed to tens of thousands of images of a multitude of "objects of interest" such as cells, bacteria, yeast, and intestinal parasite eggs, as well as the likes of crystals in the urine. All of this allows for a comprehensive, accurate identification and robust classification, within minutes[1-4] – avoiding the aspects that challenge shallow learning systems.

In the algorithm we trust

The method employed to train the AI is a logical focal point for curiosity, and the question often asked among those developing this technology is, "If the AI was a specialist, would I hire them?"

Depending on the developer, the validation process will differ. There is a generally accepted process to firstly analyse, adapt and ensure data quality (the images and their identification in this case), to compare, evaluate and develop the deep learning model (this will be a framework based on objectives and constraints), and then to test it on unseen

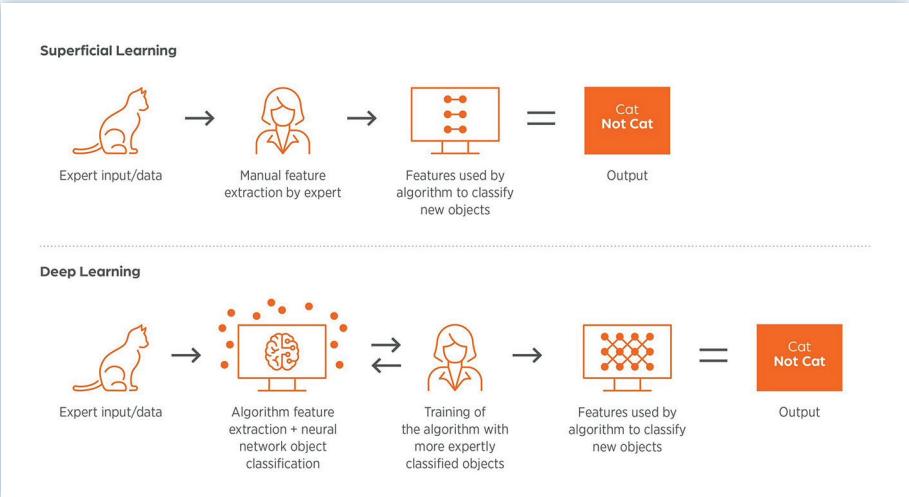

data (such as parasite eggs and non-parasite objects on faecal flotation slides).

Once it's been successfully tested across different scenarios that reflect real-world situations or challenges, it can be rolled out more widely where its intelligence increases in line with the growing size of its data library. However, it doesn't stop there; the developer should continue to monitor the performance, behaviour and impact of the model over time and in different environments to ensure that it is operating efficiently.

It's worth emphasising that the technology will never be "better" than human experts – it must first be trained by experts during that data quality phase before it can improve.

If we look to urine sediment examinations as an example, Zoetis trained its deep-learning AI (integrated within the Vetscan Imagyst's application) using thousands of iterations of urine sediment elements identified by board-certified pathologists. Those expertly classified elements were then broken down by the algorithm to the pixel level, and data from those pixel-level images and their interrelations are used by the algorithm to evaluate a sample. Algorithm validation is a rigorous process.

In the case of the AI Urine Sediment application within the Vetscan Imagyst, canine and feline urine samples were collected by Zoetis reference laboratories across the US and prepared via the platform's preparation method. Those samples were then read by boarded clinical pathologists, using manual microscopy, and scanned with the Vetscan Imagyst digital scanner. The algorithm evaluated the scanned samples, and those results were found to be comparable to the clinical pathologist's expert reads[3].

A multitude of possibilities

The evolution of deep learning AI diagnostics in veterinary medicine has already made significant progress. Beginning with faecal diagnostics, AI diagnostic capabilities have diversified to include blood smear analysis, dermatology, and now urine sediment examination as of February 2024, with ongoing advancements anticipated in the coming months.

In practice, the most noteworthy impact from the use of deep learning AI is the expedited time to diagnosis for a wide array of disease states, or highlighting where additional expert analysis or testing may be required.

The aforementioned example is only one validation process that has been implemented in the field of veterinary diagnostics. The deep learning nature of systems like this means the performance will only continue to improve.

As an added layer of assurance, in the case that a vet would like a second expert opinion – some modern systems working with AI do also provide an option to share the digital slide with a network of (human) specialist consultants or pathologists for review, who are able to respond with their assessment within hours.

Looking ahead, this approach shows how technology is shaping the future of veterinary care, bringing in new possibilities and different approaches.

That's for the future of AI to determine. Meanwhile, my waiting room isn't getting any emptier...

1. Nagamori, Y., Sedlak, R. H., DeRosa, A. *et al* (2021) Further evaluation and validation of the VETSCAN® IMAGYST™: in-clinic feline and canine faecal parasite detection system integrated with a deep learning algorithm. *Parasites Vectors* **14** (89); https://doi.org/10.1186/s13071-021-04591-y
2. Data on file, Study No. D870R-US-21-045, 2021, Zoetis Inc.
3. Data on file, Study No. DHXMZ-US-23-218, Zoetis Inc.
4. Data on file, Study No. D870R-US-22-053, 2022, Zoetis Inc.
5. Data on file, Study No. DHX6Z-US-22-166, 2022, Zoetis Inc.

Richard E. Goldstein, DVM, DACVIM, DECVIM-CA, vice president and chief medical officer for Global Diagnostics at Zoetis, is a multi-award-winning veterinarian specialising in small animal internal medicine with 12 years' experience as a professor at Cornell.