

RCVS AI Roundtable

20 May 2024

Institute of Engineering
and Technology | London





Setting the scene

Artificial intelligence (AI) is everywhere and is progressing at an ever-increasing rate. As the regulator of the veterinary professions, we have a responsibility to make sure that we are maximising the use of AI for good, but also preventing it from being used in any way that could cause harm to animal health and welfare or public health.

On Monday 20 May 2024, we held a roundtable event at the Institute of Engineering and Technology (IET) in Savoy Place, London, bringing together a range of stakeholders to help feed into our initial thinking around the regulation of the use of AI across a range of veterinary sectors.

The event, which comprised a mix of talks, plenary sessions, and small group discussions, brought together over 100 individuals from across the veterinary professions, including representatives from associations, educational institutions and technology companies, as well as veterinary surgeons, veterinary nurses and students. It also included participants from other professional regulators and the public sector, to ensure a wide range of expertise was taken into account.

The aim of the event was to help establish our initial direction of travel for the regulation of AI within the veterinary professions, to understand where the key challenges lie, and what the next steps would be in addressing these. This report serves as a summary of the day.

The aim of the event was to help establish our initial direction of travel for the regulation of AI within the veterinary professions



RCVS AI Roundtable Agenda

9.30-10.00

Registration and Coffee

Session one – background and stimulus

10.00-11.40

Welcome and introduction – Lizzie Lockett, RCVS CEO

AI and its use in clinical practice – Liz Barton MA VetMB MRCVS MCIPR, Head of Communications, VET.CT

AI and data in the veterinary space – Dr PJ Noble BSc BVM&S PhD FHEA MRCVS, Senior Lecturer in Small Animal Internal Medicine, University of Liverpool Veterinary School

AI and higher education – Dr Chris Trace BVetMed GCLT FHEA NTF FAcadMED MRCVS, Head of Digital Learning, University of Surrey

AI and regulation: managing risks – Joanna Conway, Deloitte Partner, AI & Digital Regulation, Deloitte LLP

Q&A panel

11.40-12.00

Coffee break

Session two – risks and opportunities

12.00-12.55

How are you feeling – what are the opportunities, what are the risks? – Lizzie Lockett, RCVS CEO

What would the veterinary world look like if AI worked perfectly for us? – Eleanor Ferguson, RCVS Registrar

What would the veterinary world look like if AI followed the doom narrative?

12.55-1.30

Lunch

Session three – getting down to the issues

1.30-2.45

Table discussions on the following themes (choose three):

Darren Tysoe, RCVS Chief Technology Officer

1. What does it mean to be a vet, and what can only be done by an MRCVS?
2. Should we help vet/VN students make the most of AI in learning and assessment or ban its use?
3. How do we ensure vet/VN students are best prepared to safely use AI in the workplace on qualification?
4. If AI is used in research or policy development, should this be transparent?
5. What does responsible use of AI in a clinical setting look like?
6. How should AI-led devices (inc software) used by vets be regulated?
7. If a non-MRCVS outside the UK is using AI or robots to assist a vet inside the UK and something goes wrong, who is liable?
8. What is the risk of not using AI?
9. How can veterinary professionals best work with clients/owners/keepers to ensure safe and productive use of AI?

2.45-3.10

Tea break

Session four

3.10-3.55

Feedback – Linda Belton, RCVS Junior Vice-President

3.55-4.00

Summing up, what's next, and close – Lizzie Lockett, RCVS CEO



Session one Background and stimulus

Welcome and Introduction – Lizzie Lockett, RCVS CEO



The day began with a welcome from RCVS CEO Lizzie Lockett, in which she gauged, via the Slido live polling platform, how the delegates were feeling about AI, and outlined the principles and aims for the day.

She began by asking delegates to answer a question on Slido: how are you feeling about AI now?

Answers to this question, shown via a word cloud, varied from excited and intrigued, to cautious and unsure. In response to this, Lizzie noted that this was the exact reason why the RCVS had decided to hold the event. She said that AI was something that was here to stay and something that we were all still learning about and would continue to need to learn about. It was, therefore, natural to feel trepidation. She explained that this was why the IET was the perfect place to hold the event. The purpose of the IET was to share knowledge about science, engineering, and technology, and it provided a good example of how we need to change with the times.

Lizzie then gave a brief overview of some commonly-asked questions when it comes to AI. She presented an example of what happened when she typed 'artificial intelligence will...' into a search engine, to see what came up most frequently when it came to thinking about AI on a global platform. The answers provided mainly negative connotations, with predicted endings to the 'artificial intelligence will...' sentence, including 'take over the world', 'take away jobs', 'eventually replace humanity', and 'destroy humanity' – highlighting the fact that



RCVS AI Roundtable



many people were uncertain when it came to the impact that AI could have on us all in the future.

Next, Lizzie went on to explain how it was natural to fear change and cited a few of the main things that have changed in the work environment over the course of the past 35 years. This timeline spanned from typewriters, to disk operating systems (DOS), to Windows, to dial-up modems, to palm pilots, to smartphones, and now, AI. There had been a continuum when it came to technological change, but the speed of development was increasing exponentially.

Lizzie summed up her introductory talk with an overview of why delegates had been invited to the event and what she hoped the outcomes would be. She stated that AI would not go away and that it would have a massive impact on the veterinary world and in all the areas that the RCVS regulated, from education through to professional conduct. She also said that, while the RCVS had



RCVS AI Roundtable



no remit to regulate AI and AI-powered devices themselves, the College did have the power to regulate how they were used, and that it was important to work together with the professions and innovators on how the College should go about doing this. She then spoke about the risks involved in ‘under regulating’, or ‘over regulating’, and the need not to ‘regulate people’ but to ‘regulate with people’, hence the willingness to have a large range of stakeholders involved in these vital conversations.

Lizzie said that, via the roundtable the College was aiming to understand the hopes and fears of stakeholders in a bid to put together some draft guidance on AI use, which would then be put to the Standards Committee, the RCVS group responsible for the development of guidance for the veterinary professions. She also mentioned that animal owners would be an important part of the conversation, and that the RCVS would consult its newly-formed Public Advisory Group on the use of AI, in order to take into account the animal-owning public’s perspective. She explained that, once an initial direction of travel in terms of AI regulation within the veterinary sector had been established, and when there was some more solid work to review, there would be ample opportunity for the wider professions to get involved and have their say. She added that the purpose of this initial roundtable event was to help the RCVS to understand where the key challenges lay and what the next steps would be in addressing these.

Finally, Lizzie reminded delegates that the day aimed to be interactive, fun and imaginative, and that there would be no right or wrong answers, or stupid questions. She urged delegates to enter into the discussions and listen to the talks with open minds. ●

Awareness of the risks involved in ‘under regulating’, or ‘over regulating,’ and the need not to ‘regulate people’ but to regulate ‘with people



AI and its use in clinical practice

Dr Liz Barton MA
VetMB MRCVS, Head of Communications, Vet.CT

Biography

Dr Liz Barton has enjoyed a varied clinical career, alongside diversifying through initiatives such as Vet Mums and WellVet to support those living and working in the profession. This led to career diversification, including everything from event management to website development, public speaking and content creation. She joined VET.CT as Head of Communications in 2021. The company provides remote specialist support to veterinary teams globally through technology. Both as a clinician and as part of her role, she has explored themes around the future of veterinary practice and responsible advancement of the profession, including co-producing a white paper on artificial intelligence in veterinary medicine.

What is AI?

After Lizzie's introduction, she handed over to Dr Liz Barton who spoke about AI and its use in clinical practice in the first of the expert presentations.

Liz began by exploring what is meant by AI in terms of clinical decision making and patient care. She first gave a definition of AI as "the simulation of human intelligence in machines designed to perform tasks that typically require human intelligence, such as visual perception, speech recognition, decision-making, and language translation". She then went on to define intelligence as "the ability to solve complex problems or make decisions with outcomes benefitting the actor".

To help understand the current status of AI, Liz provided an analogy, encouraging delegates to think the about human



In this “sea of human intelligence”, AI could be considered as currently comprising distinct islands, but that slowly these islands were growing



RCVS AI Roundtable



intelligence as a “sea”. In this “sea of human intelligence”, AI could be considered as currently comprising distinct islands, but that slowly these islands were growing. There were gaps between these islands of intelligence, which we did not understand, but these gaps were gradually decreasing as AI technology progressed.

She said that, currently, we were still in the world of artificial narrow intelligence, or weak AI. This included machine learning, which could be as good as, or better than, human intelligence at completing a single task.

However, the next stage, Liz explained, was a singularity, whereby the previously distinct intelligences coalesce, and the AI becomes as good as, or better than, human intelligence at completing all human tasks. This was artificial general intelligence (AGI) and could be seen as a land mass covering the sea of human intelligence. She said that, since the advent of large language models, it was predicted that AGI could become present in society as early as the 2050s. After AGI, artificial super intelligence would arrive, that was, general intelligence far superior to even the most brilliant human intelligence. However, she said it was too early for us to tell how this would present itself and what this might look like.

Liz then explained how we had come to achieve high levels of machine learning (machine learning being a subfield of AI that involves building models that can learn from data, without being explicitly programmed. Machine learning algorithms can improve their accuracy over time with more data). She spoke about the first artificial mathematical neuron, called the perceptron, which was implemented in 1958 by Frank Rosenblatt, and which closely mimicked aspects of the biological neuron.

Since then, she explained, vast neural networks of inputs, processing, and outputs have been developed to create the AI networks, also known as attention networks, that exist today. Most neural networks, Liz explained, were made up of inputs, hidden layers and/or memory/recurrent cell layers, and outputs. The number of these was only limited by the complexity of the problem, the amount of data, and the ability to train the neural network. As these networks increased in

Vast neural networks of inputs, processing, and outputs have been developed to create the AI networks, also known as attention networks, that exist today



complexity, the different neural layers began to interact with one another, reinforcing the networks and allowing them to become stronger. This was how machine learning was achieved, she explained.

She said that attention networks were based on large language models, which allows for training on sequential data to be performed in parallel, for example, applications in text generation such as ChatGPT, dictation software, and Google Translate. Liz explained that these large language models were particularly special because of their unexpected 'emergent' capabilities. For example, they could create content based on prompts and have a wide range of varied 'personalities'. However, the data that we input into these programmes may inform how political they become and whether or not they became a force for good. If we fed AI good data, then we would achieve good outputs.

Applications in clinical practice

Next, Liz spoke about applications of AI in clinical practice, focusing primarily on small animal practice. She began by highlighting the areas of impact in veterinary medicine.

In the clinical practice field, these included preventative, diagnostic, treatment and prognostic applications, as well as wider practice-based applications such as operations, research, communication, education/training and sustainability.

She then broke down these uses into further examples of potential application, which are listed below.

Clinical uses – preventative and predictive


AI could be used to screen for infectious disease or seasonal pathologies and for extracting alternative data sets. At a population level AI could be used to analyse large-scale population data, for example, around the geographic spread of a disease to predict the likelihood of disease outbreaks. On an individual animal level, AI applications included wearable data tracking and alerts, for example, movement monitors to highlight early-onset arthritis. She said that AI could also be used for genetic disease screening.

At a population level AI could be used to analyse large-scale population data, for example, around the geographic spread of a disease to predict the likelihood of disease outbreaks



Clinical uses – diagnostic

On a diagnostics level, Liz said AI could be used to analyse multifactorial data – for example, predicting the likelihood of visceral leishmaniasis/hypoadrenocorticism in dogs. One area where AI was having a big impact was on image analysis, with AI tools leading to automated measurements and adjustments in radiology, being able to predict the likelihood of the presence or absence of a specified pathology, image mapping pictures to MRI findings to look for conditions such as chiari-like malformation (CLM), and carrying out gait analysis in horses and dogs based on video inputs. In terms of cytology/histopathology, AI tools could be used to identify mitotic figures/worm egg counts and be used for pattern recognition, for example, looking at cells in the periphery in oncotype dx cancer.



Clinical uses – treatment/monitoring, prognosis

As regards ongoing treatment and monitoring of animal health and welfare, Liz said AI tools could be used to monitor treatment response, for example, the likelihood of seizure activity in dogs, and allow for continuous blood-glucose monitoring. In future it could allow for individualised medicine with tailored case management taking into account a variety of factors specific to the animal and client, including clinical history, geography, genetics, other medications and allergies.

In terms of prognostics, Liz highlighted that AI could be very helpful when making high-stakes decisions, but highlighted that tools needed to be clear on their accuracy and limitations, especially when they fed into life and death decisions. An example of this would be in predicting surgical outcomes, for example, from colic in horses or acute abdomen in dogs.

Non-clinical - operational

Non-clinical AI applications could benefit practices around increased efficiencies and enable more time to do things veterinary professionals love and were good at, instead of admin. For example: automated transcription/translation, clinical notes, owner communications, content creation for websites and social media, triage and case prioritisation and improving time management through smart scheduling. ●

Non-clinical AI applications could benefit practices around increased efficiencies and enable more time to do things veterinary professionals love and were good at



AI and data in the veterinary space

**Dr PJ Noble BSc BVM&S
PhD FHEA MRCVS, Senior
Lecturer in Small Animal
Internal Medicine, University
of Liverpool Veterinary School**



Biography

Dr PJ Noble works in the Small Animal Teaching Hospital at the University of Liverpool's Leahurst Campus as part of the internal medicine team. Here he oversees internal medicine referral cases and trains undergraduate students in small animal internal medicine. With nearly 30 years of experience, he enjoys encouraging students to develop the complex, integrative thought process required to manage these cases.

Dr PJ Noble is also an investigator on the Small Animal Surveillance Network (SAVSNET) project. SAVSNET collates a growing database with over 11 million clinical notes from over 500 veterinary premises across the UK. His key interest here relates to text-mining clinical notes to gain insights into companion animal health. This brings together interests in small animal medicine and computer programming. PJ has written a suite of tools (SAVSNET Datalab) used by undergraduate and post-graduate students to pursue text-mining projects using SAVSNET data. Through collaborations with the Departments of Computer Science at Liverpool and Manchester and Durham University, he also pursues the use of the latest methodologies in machine learning/AI, primarily using neural language models to unlock the potential of the SAVSNET database.

Presentation

PJ opened his talk by explaining the role of the SAVSNET project. He explained that SAVSNET harnesses electronic health

The importance of carrying out processes and using data in an ethical manner, to ensure that data was not shared in inappropriate ways, as well as making sure that data handling was secure



and environmental data for rapid and actionable research and surveillance. SAVSNET's research priorities include antimicrobial use and resistance, climate and environment, and infection and zoonosis.

Its key aims, he explained, were to:

- monitor disease trends over time and highlight appropriate interventions;
- identify populations at risk and monitor treatments and outcomes;
- provide data resources for academics and others;
- improve general public awareness of small animal diseases and prevention; and,
- provide a route to clinical benchmarking for vets in small animal practice.

SAVSNET's angle on data

PJ said that SAVSNET aimed to gain as much data as possible and turn it into actionable intelligence. The data gained by SAVSNET comprised pet data and clinical notes. PJ noted that at first, the team did not think that the free-text data from clinical notes would be useful, as it could be difficult to assess and annotate. He noted that large language models could get confused, but that they were keen to use them to create accurate information for each use case, as well as explainable results that would be able to show traceback in order to check that the information had viable applications. He also emphasised the importance of carrying out processes and using data in an ethical manner, to ensure that data was not shared in inappropriate ways, as well as making sure that data handling was secure.

PJ then went on to explain the use of large language models, and the fact that their use was starting to increase (ChatGPT being an example of a large language model). Large language models were AI systems that processed vast amounts of text data to understand and generate human language. They used a technique called unsupervised learning, where the model was trained on a large amount of data without any specific labels or targets. The goal was to learn the underlying structure of the data and use it to generate new data that was similar in structure to the original data. Large language models were designed to capture complex relationships between words and phrases.

Large language models were AI systems that processed vast amounts of text data to understand and generate human language



PJ explained that the development of effective large language models required access to extensive datasets, which was particularly evident in the field of veterinary science, where datasets were continually expanding. These growing datasets needed to be integrated with advanced computing capabilities, and both must advance at the same rate. However, it was important to note that these models also learned from the biases present in the data. This raised concerns about the possibility of the models generating inappropriate and biased outputs. Sometimes models would also 'hallucinate', creating false outputs. To mitigate this, Dr Noble said it was essential to fine tune the models using specific tasks designed to filter out bad language. However, the question of who determined what constituted an appropriate output remained critical.

PetBERT

PJ explained the development of SAVSNET's PetBERT, a masked language model, trained on over 500 million additional words from first-opinion veterinary clinicians from across the UK. He explained how the tool had been developed from the BERT model of machine learning but made bespoke to the veterinary profession through the inputting of veterinary clinical notes. This allowed PetBERT to produce improved outputs in its specific role in detecting gastrointestinal (GI) disease in dogs. PJ said that PetBERT could classify records and tell whether or not GI disease was present in the patient. The model was shown to be able to recognise the presence of GI disease in a patient two weeks earlier than usual.

In terms of explainability, PJ said that large language models were difficult, and it could be challenging to retrace the outputs to the original inputs. However, it was possible to ask why it gave you the answer that it did.

Ethics

In thinking about the ethics surrounding the building and use of large language models, PJ noted that it was important to ask the following questions:

- How do you get your data?
- How do you curate your data?
- How do you generate signals accurately and securely?

The development of effective large language models required access to extensive datasets, which was particularly evident in the field of veterinary science, where datasets were continually expanding



RCVS AI Roundtable

- How do you use and explain the use of outputs?
- How do you avoid bias?

He noted that the most important thing, when it came to ethics, was being able to understand how your AI system worked so that you could explain how the outputs were arrived at.

Conclusion

PJ concluded his talk by stating that AI development, both software and hardware, was moving fast – it was expanding and advancing all the time and things were constantly being updated. He noted that, in a veterinary context, the use of large language models could save masses of time and could allow huge studies to be carried out in a faster, more accessible, and more cost-effective manner. He highlighted that different models could be used for different tasks and that outputs could be made explainable, be tested for accuracy, and be used securely. Finally, PJ reinforced the fact that ethical use of large language models required an understanding of the AI by the user. ●



Different models could be used for different tasks and that outputs could be made explainable, be tested for accuracy, and be used securely



AI and higher education

**Dr Chris Trace BVetMed
GCLT FHEA NTF FAcadMED
MRCVS, Head of Digital Learning, University of Surrey**


Biography

Dr Chris Trace is the Head of Digital Learning within the Surrey Institute of Education. He is responsible for leading the Digital Learning Team in developing, implementing, and supporting academic staff to use contemporary technological systems, tools and practices in ways that enable successful delivery of the education strategy. Chris initially joined Surrey in 2013 as the School of Veterinary Medicine's eLearning Coordinator, with a remit to support the use of technology to enhance the curriculum. Whilst at the vet school, Chris also took on the roles of Undergraduate Admissions Tutor, Pedagogical Development Coordinator, Chair of the Student Staff Liaison Committee and Faculty Theme Champion for Scholarship of Teaching and Learning.

Presentation

Chris opened his talk by saying that his institution was taking a leading role in operationalising AI in the educational sector and that it was currently working with the Department of Education in training secondary school students in AI skills.

Focusing on the role of generative AI in the university sector, he said that the appearance of OpenAI's ChatGPT on the scene in November 2022 led to a boom in the sector – with thousands of generative AI tools now available. He added that, within this relatively short period of time, ChatGPT and similar tools had already managed to 'smash through' some of their initial limitations and that their capabilities would only grow, saying that "whatever we think a platform can do now, it is at its least capable point".



The appearance of OpenAI's ChatGPT on the scene in November 2022 led to a boom in the sector – with thousands of generative AI tools now available



RCVS AI Roundtable



He said the most important word in caveating what generative AI could not currently do was 'yet'. He then went through how generative AI had already started to be used in several different areas of higher education. In terms of assessment and feedback, Chris said it had helped make improvements in academic integrity, had led to the rise of AI-assisted marking and led to increased expectations from students around assessment and feedback. He pointed to tools such as the Surrey-developed Keath.ai (Key Evaluation and Assessment Tools Hub), an assessment system which provided personalised feedback and a full justification report for a number of different 'inputs,' including essays, dissertations and case studies.

We still have a choice and control on whether AI would lead to a eutopia, or a dystopia

In terms of teaching, Chris pointed out that generative AI had had an impact on lesson planning, developing interactive scenarios and developing critical thinking skills. Again, he highlighted a tool in use at Surrey called 'Virtual Reality in Veterinary Practice,' which uses inputted clinical data to develop a series of interactive clinical veterinary scenarios – which include virtual avatars – for students to use in their learning.

From the student's perspective, Chris said generative AI tools were helping to override digital inequalities amongst students, allowed students to test and consolidate their knowledge, explore key topics in greater depth and develop their prompt engineering skills so that they were properly structuring the instructions to generative AI tools to get the optimal results. A tool in use at Surrey in this respect was Tutello, an AI chatbot into which students could feed module content and then train it to be able to contextually answer any questions they may have about their course.

In judging the merits and demerits of AI in higher education, Chris shared with the audience a SWOT analysis (strengths, weaknesses, opportunities, threats) that the University of Surrey had conducted. Weaknesses included lack of emotional intelligence and intrinsic bias; identified threats were overreliance on AI tools and job displacement; while opportunities included enhanced learning and cost savings; and strengths included efficiency, availability and scalability.

Chris concluded by saying that, at the stage when generative AI was still dependent on human inputs and had not reached self-taught general intelligence, we still had a choice and control on whether AI would lead to a eutopia, or a dystopia. ●



AI and regulation: managing risks

Joanna Conway, *Deloitte Partner, AI & Digital Regulation, Deloitte LLP*

Biography

Joanna is an AI and digital regulation specialist and partner in Deloitte's global multidisciplinary digital regulation practice. She advises global technology clients on responding to digital and AI regulation, supporting as a consultant on regulatory governance, risk and compliance. As well as her consultancy work, Joanna is also a practising digital and regulatory lawyer advising on emerging technology and AI laws and risks. She sits on the Law Society's Intellectual Property Committee and she and her team won the coveted FT Europe Legal Award for their digital regulation work.

Presentation

Joanna, a partner at Deloitte specialising in AI and digital regulation, said her talk would focus on AI in the context of ethics and regulatory compliance

She said that, as part of her role with Deloitte, she worked with clients on putting AI into practice. She said AI was an adaptive and, to some extent, autonomous technology that was best used in making or aiding decisions but that its outputs were always based on the inputted information.

She added that AI could not be looked at as a technological innovation in isolation, saying that a lot of the work carried out by Deloitte was about looking at other technologies that may be coming – for example, quantum computing and the metaverse – and considering how these could be best harnessed not only in terms of efficiency and cost savings, but also in terms of safety, ethics and risk mitigation.



AI was an adaptive and, to some extent, autonomous technology that was best used in making or aiding decisions but that its outputs were always based on the inputted information




Moving onto the subject of risk, she asked delegates to consider how broad the spectrum of risks might be. She said that there was already a lot of law and regulation that, while not being specific to AI, impinges on its use, for example, around the use of personal data, which she said was something that held businesses and organisations back from using AI products. She added that there were risks in terms of use and misuse of intellectual property, with many large language model developers currently facing litigation over the provenance of the data they input and the fact that the outputs may also breach intellectual property rights.

In terms of ethics, she said that there had been a huge amount of regulation coming out of the European Union, with the EU AI Act likely to have a ripple effect in a number of other jurisdictions, including the UK. However, she added that it was not just about legal compliance, and that businesses and organisations needed to consider where the use of AI fitted in with their corporate and social values, for example, in terms of the potential for workforce displacement.

Joanna said that, because the core principle of AI use needed to be that we should dictate to the AI and not the other way around, it was useful for all organisations to have a corporate AI framework to ensure that they were being consistent in terms of when and why AI tools were used, as well as being consistent with corporate values. Her advice was that the foundations of any corporate AI framework should be risk management, considering all the different risks in AI use and mitigations to manage these identified risks. To ensure that an AI framework worked in practice there needed to be guidance and policies for those in the veterinary sector and that companies and organisations needed to upskill and train people in getting the most use out of AI.

Joanna's overall advice on AI use was that, while it could do much to cut costs and improve the efficiency and quality of work, it didn't need to be used for absolutely everything. The key applications for AI were for problem solving and making things better but that it needed a strong risk management framework underneath it, and should be guided by strategic aims around what companies and organisations wanted to get out of AI use. ●



Companies and organisations need to upskill and train people in getting the most use out of AI



Q&A Panel



Following each of the individual presentations, RCVS CEO Lizzie Lockett chaired a panel discussion picking up some of the key discussion points from each speaker and inviting delegates to ask questions from the floor.

How to harness AI

The first question was around how companies and organisations could best harness AI, with Joanna Conway saying that they needed expertise and knowledge in AI to inform any organisational strategy they may have and to take a centralised approach to its appropriate use. This way organisations would avoid a patchwork and inconsistent policy on AI use. Dr PJ Noble added that organisations should have an agile system for disseminating AI innovation and appropriate regulation, using it to save time on routine tasks and using that spare time to be more creative.



AI in higher education

The next question concerned the use of AI in higher education and whether we can use AI to make teaching more uniform. Dr Chris Trace said that the Russell Group of universities recently published five principles on the use of AI in education, which formed a good foundation for its use and included items about AI literacy and thoughtfully embracing new technologies. He said that Jisc, a not-for-profit that provided network and IT services and digital resources in support of further and higher education and research, and the Quality Assurance Agency for Higher Education (QAA), were also taking a lead on the issue within the sector. Both organisations were looking at what other industries were doing and working with heads of e-learning at universities to find out what was working and what more could be done.



He added that some universities liked their individuality in terms of how they focused their curricula and teaching, but if the sector shared intel then they could free up time to allow time for greater creativity in teaching. He added that in France all universities taught the same veterinary curriculum and that there was a shared app that used the same information.

On the same question, PJ added that universities in the UK would always want to have their individual unique selling points and it was likely that AI, rather than standardising everything, would be used to create a unique experience in each of the vet schools by freeing up time for greater creativity.




Resistance to using AI

The panel was next asked whether resistance to using AI technology was a problem in itself. Dr Liz Barton answered that in certain areas, for example, radiologists working in the human healthcare sector, refusing to use AI technology could potentially be negligent because of the enhanced diagnostic capabilities it confers. She added that those who did not adopt the technologies would just find themselves being outcompeted commercially and would need to adopt and adapt to succeed.



Who's responsible when AI goes wrong

The liability for when AI use went wrong was the next topic the panel was asked to discuss. Joanna said there was currently no certainty around this – as UK common law was iteratively developed via case law and the technology was so new. However, she said that this was where new legislation and regulation may need to step in to find some of the answers. She said that in the EU, alongside the AI Act, there were other measures looking at non-contractual liabilities. She added that if companies and organisations were thinking of deploying AI then it would be important to manage it contractually and to get consent from pet owners for any data use. She said that a risk-based approach to AI use was the best way forward and that therefore AI shouldn't be used to make decisions, but should be used as an aid in making decisions.



On the same question Liz spoke about the importance of transparency with animal owners over the use of their pets' data, while PJ added that one of the difficulties for risk and liability in the AI field was understanding the huge range of activities and projects that the technology might be involved in and all the associated risks. He added that the form of words the EU had arrived at in its legislation was useful as it was more about the effects of its use – for example, on the labour market – but that there was still a massive regulatory challenge.



Helping pet owners understand AI

The final question for the panel was on how to help pet owners understand the differences between different AI tools in terms of their quality. Chris said it was about learning to talk to clients about the potential risks of using low-quality AI tools and 'Dr Google' especially if it was giving poor or dangerous advice. ●



Session two

Risks and opportunities

The next part of the event focused on the opportunities and risks when it comes to AI in a veterinary setting.

This session was opened by Lizzie Lockett, who encouraged delegates to form 'people clouds'. Delegates were asked what they thought the biggest risk was in terms of AI use in veterinary medicine, then those who agreed with that particular person's perspective were asked to go and stand alongside them. The point of the exercise was to gain a visual representation of what people believed the biggest risk to be in terms of veterinary AI. Risks raised included the loss of humanity, that is, the loss of the human element in client-patient interactions, as well as legal liabilities, with questions around who should take responsibility for the use of these tools – should it be the people creating the devices, or those using them?

A further risk cited was the risk of clients bypassing veterinary surgeons and going straight to AI tools for formal diagnoses and treatment without properly understanding how the AI worked and why it was giving them the advice it was. The data inputted into that particular AI programme may not have been high quality, and even if it was, there was no saying that it would have reached a correct outcome. It all relied on the quality of the data, and being able to trace the data back to a reputable source. The people using the AI needed to understand the complexities of the data being put in and how the system worked.

This led onto the biggest risk, which was cited as 'bad' data being entered into AI systems, resulting in poor or false outcomes – AI relied on the quality of the data being inputted. If the data being inputted was of bad quality or incorrect, then the AI would be likely to form biases and generate further false





RCVS AI Roundtable



and inaccurate information, which would, inevitably, have a negative impact on patient outcomes.

RCVS Registrar, Eleanor Ferguson, then asked delegates to work in groups to discuss and illustrate what they thought the veterinary world would look like if AI worked perfectly for the veterinary world, and the potential 'doom' scenario if it did not – ie utopia and dystopia.

The utopian view generally centred on the idea that AI would be used to assist and enhance veterinary medicine by providing advice to both practitioners and clients. However, this view also contended that veterinary medicine should ultimately be human-led and the AI should be used as a tool to support upholding animal health and welfare, not as a replacement for veterinary professionals. It was also suggested that AI could be used to free up the time of vet teams, to help automate administrative processes to give teams more time to work on more important clinical matters. It was suggested that AI could be used to solve issues such as anti-microbial resistance, cure infectious diseases, and support research through analysing large data



RCVS AI Roundtable



sets to speed up research processes. It was suggested that AI could also result in more affordable treatment for animals and help support contextualised care.

In terms of dystopian views, these centred mainly on the risk that AI could take over, rendering the vet irrelevant, and compromising animal health and welfare. Other scenarios cited were that it could create a monopoly by becoming primarily focused and based on single pharmaceutical companies. It could also lead to clients trusting AI over vets, vets becoming too reliant on the technology, and AI reinforcing bias through inappropriate biased programming.

Please note, the above utopian and dystopian discussions were designed to get delegates thinking about the best possible and worst possible scenarios in the far future of veterinary AI. There was an array of opinions expressed and the above summary is designed to give a brief flavour of the discussions. ●



Session two

Getting down to the issues

The afternoon session was led by Darren Tysoe, RCVS Chief Digital Officer, with the room being split between nine tables, with each table having a particular question to consider. These questions, listed below, were about the immediate practicalities of what guidelines and guardrails could be put in place for the appropriate and risk-conscious use of AI in the veterinary context. Over the course of an hour and a quarter, each delegate was asked to move around the tables, to consider at least three of the questions, which were:



- 1. What does it mean to be a vet, and what can only be done by an MRCVS?**
- 2. Should we help vet/VN students make the most of AI in learning and assessment or ban its use?**
- 3. How do we ensure vet/VN students are best prepared to safely use AI in the workplace on qualification?**
- 4. If AI is used in research or policy development, should this be transparent?**
- 5. What does responsible use of AI in a clinical setting look like?**
- 6. How should AI-led devices (including software) used by vets be regulated?**
- 7. If a non-MRCVS outside the UK is using AI or robots to assist a vet inside the UK and something goes wrong, who is liable?**
- 8. What is the risk of not using AI?**



9. How can veterinary professionals best work with clients/owners/keepers to ensure safe and productive use of AI?

Following a short break after the discussions, each table was asked to report back on the discussion outcomes and some of the conclusions that were reached. This session was chaired by RCVS Junior Vice-President Linda Belton.

What does it mean to be a vet and what can only be done by an MRCVS?

Reporting back on this question was Professor Tim Parkin, Head of Bristol Veterinary School and a member of RCVS Council. He said that the discussions were around the diagnostic possibilities of AI tools and the way that they could be used to assess and deliver information for veterinary professionals that they could then use in making their clinical, professional decisions. He said there were questions and concerns around the validity and efficacy of different AI tools and their use in practice, as well as the potential for strain on the vet-client relationship if animal owners had the choice to either go to a 'human vet' or a 'machine vet' for treating their animal. However, this was why the final decision-making process had to still lie with the vet, who would have the ultimate responsibility and liability regarding the treatment of the individual animal.

Should we help vet/VN students make the most of AI in learning and assessment or ban its use?

Reporting back on this question was Jordan Nicholls, RCVS Lead for Undergraduate Education. He said the group discussions confirmed that students were already using AI tools to support their study and so it was a question of how best to train them to use these tools ethically, transparently and appropriately rather than banning their use. Appropriate use was about using AI tools to help present their knowledge and for demonstrating their understanding of what has been produced. However, AI tools could not be used in cases where critical thinking or emotional intelligence was required, for example, in providing contextualised care. He said that there was agreement that a set of principles around AI needed to be developed to police its use but that, overall, AI tools





provided an opportunity for individualised learning plans and opening up the profession to wider groups of people.

How do we ensure vet/VN students are best prepared to safely use AI in the workplace on qualification?


RCVS Council member and past President Dr Kate Richards reported back on this question. She reported that there was a consensus that principles on the safe use of AI in clinical practice should be woven into inductions at different practices – whether this was on an intra-mural rotations (IMR)/extra-mural studies (EMS) student placement or as a paid clinical role. Allied to this would be the need for the whole practice environment to be regulated rather than just individual vets and vet nurses, although individual professionals should be told that they have ultimate responsibility for any decision-making they make based on outputs from an AI tool. She said that it should be made clear to students that, similarly to how they learn to critically evaluate academic papers, there should also be training on how to evaluate AI tools so that students and members of the professions could make the decision as to whether they were comfortable using certain tools and potentially sacrificing their professional autonomy to them.





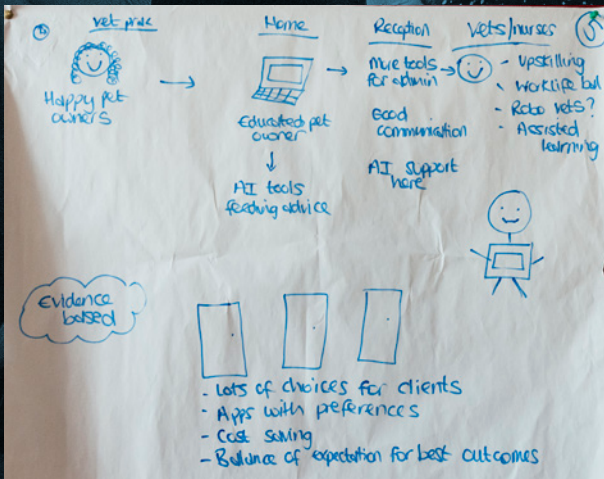

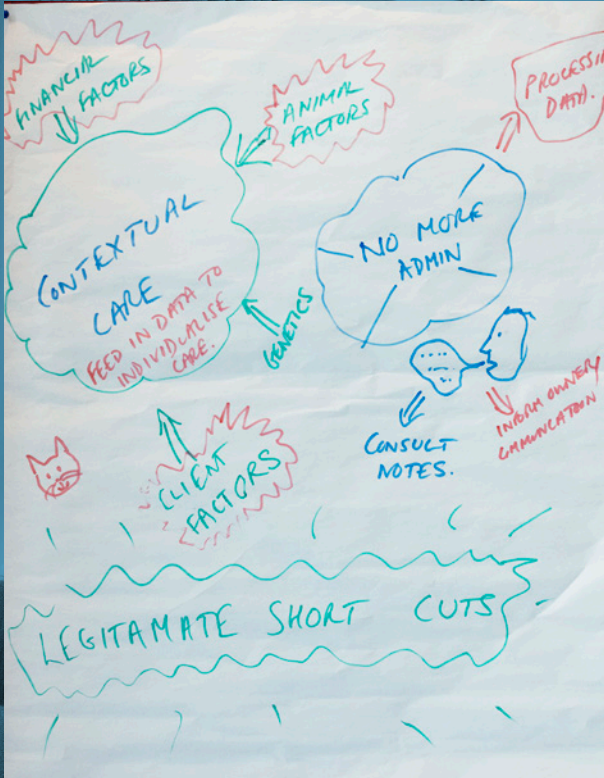
If AI is used in research or policy development should this be transparent?

Dr Melissa Donald, RCVS Senior Vice-President, spoke about some of the conclusions that her groups reached on this question. She said that the use of AI tools in research should always be declared and justified. In terms of policy decisions, for example, amongst public bodies such as the Veterinary Medicines Directorate (VMD) and the Department for Environment, Food and Rural Affairs (Defra), the use of AI tools in providing data should again be transparent. Furthermore, it was considered vital that a 'human' who could critically evaluate and ask questions of the evidence with which they had been presented had the final sign-off of any policy decision. She said there was an emphasis in discussions on the importance of transparency and trust and that the small percentage of 'bad actors' who misused AI did not ruin it for everyone else.





RCVS AI Roundtable





What does responsible use of AI in a clinical setting look like?


RCVS Head of Standards Gemma Kingswell gave a summary of the discussions on this question. She said that the main topic of debate was around transparency and consent – with one of the questions posited being that veterinary surgeons and veterinary nurses were already using various devices and machines in clinical practice without having gained specific consent for their use from animal owners – so was using AI fundamentally different and, if so, how? Her groups discussed that there could be a more tiered approach to gaining specific consent for the use of AI tools depending on the level of risk. So, the use of AI for low-risk activities may not need consent, but where clinical decision making was impacted, consent may be required – similar to the cascade in prescribing medicines. Another theme was around validation and AI and whether or not it was reasonable to expect veterinary professionals to be able to qualitatively evaluate those AI tools and devices or whether there needed to be a quality kitemark or regulatory regime similar to that of the Medical and Healthcare Products Regulatory Agency (MHRA) in human healthcare. Overall, it was considered that the veterinary professional needed to remain ultimately responsible for decision-making and that there remained a connection between the vet, tool and output to ensure that clinical decision-making wasn't wholly delegated.



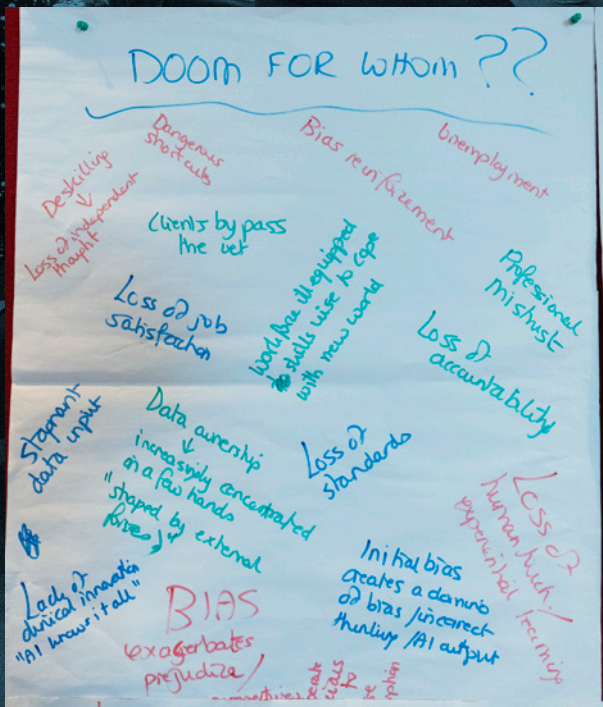


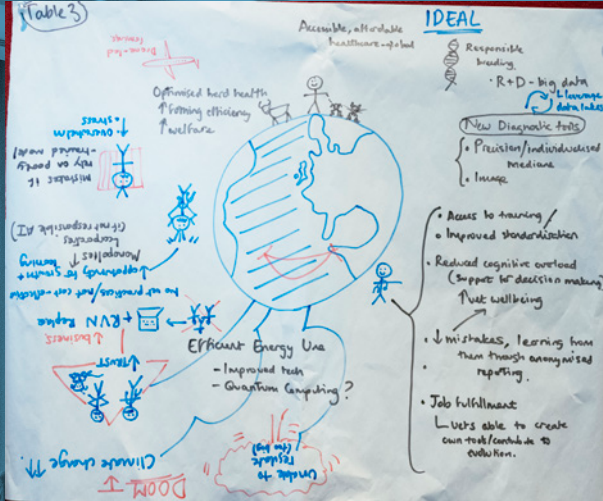
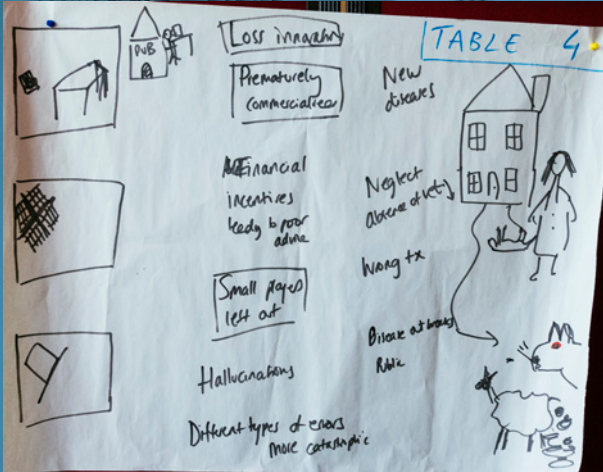
How should AI-led devices (including software) used by vets be regulated?

In answering this question, RCVS Chief Technology Officer Darren Tysoe said that, amongst his groups, there had been a general consensus that the veterinary professional should take responsibility for any outcomes that resulted from the use of an AI device or software, even if it turned out it was faulty. He said there was an agreement that there should be regulation of AI-led software and devices but that this was potentially complicated by the international nature of the products and also the fact that, despite rigorous user testing, some problems may not be apparent until further down the line. He said there was some anxiety over the matter of liability for when things went wrong, but that ultimately the veterinary professional was liable as the person who combined the insights and data from AI-led devices with their own expertise and clinical judgement to arrive at a decision.





RCVS AI Roundtable





If a non-MRCVS outside the UK is using AI or robots to assist a vet inside the UK and something goes wrong – who is liable?


RCVS Director of Advancement of the Professions Angharad Belcher said that the consensus amongst the groups was that if the animal is under the care of the UK vet, then it was ultimately their responsibility. However, in terms of other forms of liability, that may depend on their having done due diligence upon signing contracts with external suppliers. There were also discussions around quality kitemarks for AI-led products, whether the VMD's remit could be expanded to encompass AI-led devices and the need to upskill people in practice settings so they could do due diligence on contracts. The discussions agreed that there was a need for veterinary professionals to carefully consider the datasets on which the AI tools they use might depend – so, considering where the data came from, any potential inherent biases and how often the data was reviewed, in addition to considerations around data protection. She added that transparency around telling clients about the kinds of AI-led tools and devices that were used could help educate owners and enhance the relationship and trust by helping them understand why certain tools were used, or not.



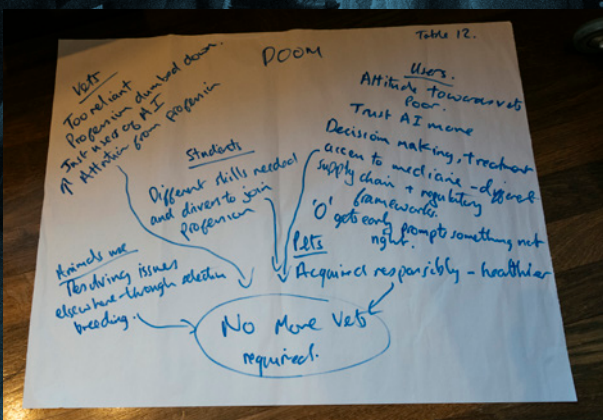
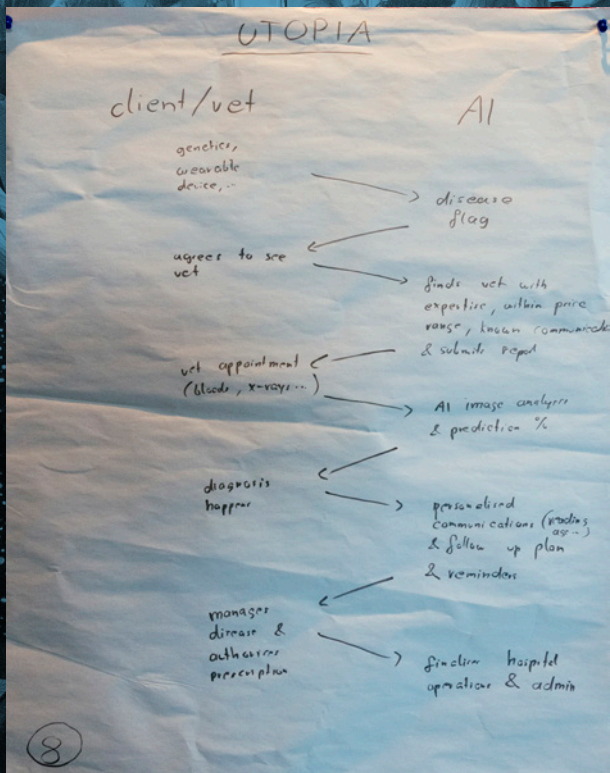

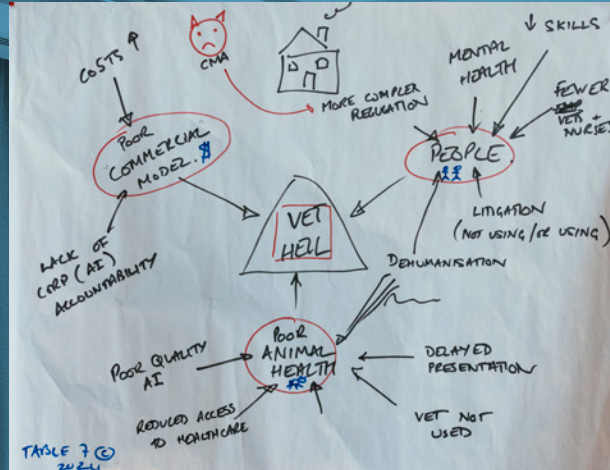
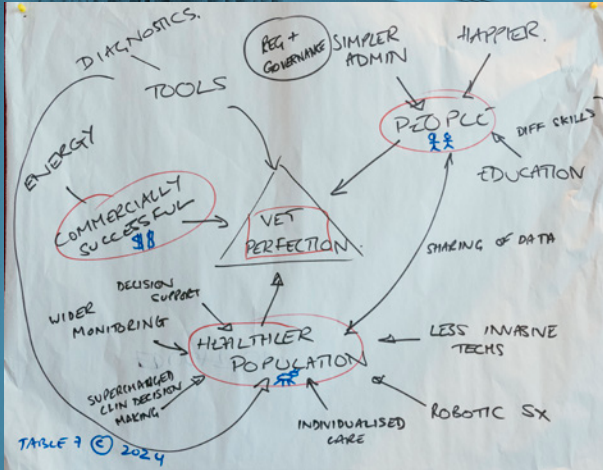
What is the risk of not using AI?

Answering this question, RCVS Registrar Eleanor Ferguson said the concerns within her discussion groups were around the fact that if veterinary professionals, individually and collectively, did not use AI tools responsibly, then someone else would step into the space and potentially make the professions obsolete. She said that there may also be an expectation amongst clients, students, and society at large that the veterinary professionals would be using AI tools if they could lead to better clinical outcomes. She said that the groups had felt there was the potential it could become negligence, and even a disciplinary issue, if veterinary professionals were not using AI tools that were known to be beneficial to patient outcomes. As to what might be missed if AI tools were not embraced by the professions, she said there were potential positives that could be squandered such as developing medicines, allowing more time to invest in quality improvement, giving more choice for clients via contextualised care, disease prevention and surveillance and enhancing





RCVS AI Roundtable





RCVS AI Roundtable



professional wellbeing. She said there was the fairly unanimous feeling that not embracing this technology would make the professions less competitive.

How can vets best work with clients/animal owners/keepers to ensure safe and productive use of AI?

Dr Louise Allum, who chairs the RCVS Public Advisory Group, said that in discussing this question there was a view that AI could be used to help members of the public make better decisions on the ownership of animals, so that the pets they own were better matched to their circumstances, for example, around cost and potential long-term health conditions. She added that there were questions for the profession around whose problems were being solved by AI and that veterinary professionals needed to think about how the tools could enhance outcomes and experiences for clients and their animals, and not just how they could make things easier for the profession. Finally, Louise said there were questions around the costs of using AI tools and that, while transparency over their use was important, would the public be happy to pay additional costs for their use in clinical practice? ●





Conclusions

In concluding the day, CEO Lizzie Lockett thanked all those present for their open and collaborative approach, and acknowledged that AI had the potential to touch on every aspect of the veterinary professions, from the way students were educated, to the way in which care was delivered to the UK's animals.

She asked people the same question with which she had started the day – 'How are you feeling about AI now?' The answers had moved on to a greater focus on the potential of AI, with the words 'excited', 'curious' and 'intrigued' being some of the most commonly-cited, which was a positive development.

Lizzie summarised the feeling in the room that any action taken by the College to regulate veterinary professionals' use of AI would need to emphasise the fact that the veterinary professional remained ultimately responsible for clinical decision-making and the delivery of care, and that apart from very low-risk scenarios, there should always be a 'human in the loop'.





RCVS AI Roundtable




She also recognised that work would need to be carried out with providers of AI-based tools to enable vets and nurses to understand the basis on which these products worked so that they could be used confidently.

Lizzie said that the next steps were for the RCVS to identify specific actions to be looked into by the relevant RCVS Committees and Working Groups. For example, the Standards Committee to consider changes to the *Codes of Professional Conduct*; the Education Committee to consider any relevant changes to vet school accreditation standards and/or Day One Competences for new graduates; the Advancement of the Professions Committee to review the need for training and culture change around the adoption of new technologies; and the Public Advisory Committee to better understand any concerns the public may have around AI use in a clinical setting. Depending on the nature and extent of any changes proposed by these committees, they may go out to consultation – this would be for consideration.

Meanwhile, an internal team at the College would develop policy and guidance for how AI may be used as part of the RCVS's day to day work. ●

The next steps are for the RCVS to identify specific actions to be looked into by the relevant RCVS Committees and Working Groups



RCVS AI Roundtable

Royal College of Veterinary Surgeons

3 Waterhouse Square

138–142 Holborn

London EC1N 2SW

www.rcvs.org.uk

info@rcvs.org.uk