

YOU CAN (AND REALLY SHOULD) BUILD YOUR RESPONSIBLE AI GOVERNANCE POLICY FROM SCRATCH. RIGHT NOW.

BY PETRA HARMS, DVM

Veterinary professionals take responsibility seriously. Our careers hinge on having a license that tells the world we can be trusted with the medical care of animals. We work hard to safeguard the privilege of having it. It's not surprising that AI technology makes some veterinarians uneasy. How can we expect to split the workload of practicing medicine with an AI tool when we carry the liability for errors?

Knowing all this, the industry's erstwhile silence on the topic of AI governance has been surprising. However, regulatory bodies and associations are finally stirring, and AI position statements are starting to trickle out ([RESOURCE 1](#)).

While the new position statements provide some suggestions for appropriate AI use, it's left up to us to figure out how to apply that framework to day-to-day practice. Establishing a governance structure is nothing new. Think about how you approach the ordering, management, administration and prescription of drugs in your practice. Medications share much in common with AI technology. They can help, but also hurt. Sometimes, they do things that are unanticipated or outside of your control. They need careful management, oversight, communication and staff training to ensure appropriate handling. That's why we have a responsible drug governance system baked into the core of our day-to-day practice.

We can do the same for AI tools. Responsible

AI governance is the principle that organizations developing or using AI tools should assess their trustworthiness, evaluate potential risks or harms and take proactive steps to mitigate those harms.

Steps to Design Your AI Governance Policy

Catalog your AI tools

Start by cataloging which AI tools you are already using or are planning to use. Include veterinary-specific tools like scribes, diagnostic aids, scheduling aids, client chatbots, medical note summarizers and generative AI for discharge notes. Don't forget nonspecific tools like ChatGPT if members of your team are using it for any reason. Keep in mind that your diagnostic laboratory may also have embedded AI technology.

Know the signs of trustworthy AI and identify the harms

To make a responsible AI governance policy, you need to know what trustworthy AI looks like. Luckily, the work has been done for us. The Organisation for Economic Co-operation and Development (OECD) outlines five principles that are characteristic of trustworthy AI systems.⁴ Next, evaluate how well your AI tools follow the principles of trustworthy AI ([RESOURCE 2](#) and

CHECKLIST). Where they don't follow the principles, you need to identify what kind of risk they may pose to your patients, staff and clients. Examples of risks include:

- ▣ Bias in training data that can cause errors in treatment decisions leading to patient harm
- ▣ Data privacy violations for clients and staff
- ▣ Errors caused by AI tool use possibly opening up veterinarians to litigation or board complaints
- ▣ Veterinarians developing an over-dependence on diagnostic aids, leading to an atrophy of their diagnostic skills
- ▣ Support staff feeling anxious about losing their jobs to AI

Assess risk levels

For each potential harm that you can identify in an AI tool, you will want to predict the likelihood of the harm occurring and the severity of the harm if it does happen. The International Association of Privacy Practitioners uses a 3 × 3 probability/severity harms matrix as an example of how to evaluate the risk level of AI tools (**TABLE 1**).⁵ We can apply this example to veterinary medicine. Each of the three categories of Severity of Harm and each of the three categories of Probability of Harm are assigned a numerical value. Then, each risk of an AI tool is evaluated based on how likely the harm is to occur (from improbable to probable) and on how severe the harm is likely to be. Marginal harm may only cause mild inconvenience, whereas critical harm can mean significant injury to a patient, staff member or hospital. Remember, the probability and severity of harm may change depending on the user of the technology — for example, a new graduate as opposed to an experienced veterinarian. Multiplying the severity number by the probability number gives you an idea of whether the AI product is high risk, medium risk or low risk.

TABLE 1 3 × 3 Probability and Severity Harms Matrix⁵

PROBABILITY OF HARM	SEVERITY OF HARM		
	CRITICAL: 3	MODERATE: 2	MARGINAL: 1
Probable: 3	High (9)	High (6)	Medium (3)
Occasional: 2	High (6)	Medium (4)	Low (2)
Improbable: 1	Medium (3)	Low (2)	Low (1)

Prioritize and mitigate risks

Once you have assigned a risk level to each AI tool, you can decide how to minimize the risks. Very high-risk AI technology (high probability combined with high severity of harm) will likely be tools that you will avoid using entirely. Moderately high-and medium-risk tools may be worth using if they bring significant benefit that can't be achieved in a different way and if you can implement ways to decrease the risk profile to a level

1

Resource: Regulatory AI Position Statements in Veterinary Medicine



American Association for Veterinary State Boards:

This white paper highlights that licensees must ensure patient care, client privacy, informed consent and compliance with legal and regulatory standards, outlining risks and calling for member organizations to provide further guidance.¹



The Canadian Veterinary Medical Association:

This position statement advocates for scientifically rigorous development and regulation of AI technologies in veterinary medicine to ensure patient safety, reduce liability and align with national standards and regional regulatory policies.²



The American College of Veterinary Radiology and European College of Veterinary Diagnostic Imaging:

This position statement emphasizes that AI in veterinary diagnostic imaging must follow rigorous standards for transparency, validation, data security and expert oversight, advocating for evidence-based development, professional education, informed client communication and regulatory guidance to ensure AI enhances rather than compromises veterinary care.³

Resource: OECD's Characteristics of Trustworthy AI and How They Apply in Veterinary Clinics⁵

1 Inclusive growth, sustainable development and well-being

From a veterinary perspective: AI tools in a veterinary clinic should support the well-being of animals, owners and staff, while reducing environmental impact — such as optimizing resource use or minimizing waste. They should also be accessible to all pet owners, regardless of socioeconomic background, language or location. They should enhance animal health and client satisfaction while contributing positively to staff well-being.

2 Human rights and democratic values, including fairness and privacy

From a veterinary perspective: AI tools must respect privacy laws and maintain patient and client privacy in line with the client privacy responsibilities of veterinarians. They must ensure fair treatment of all clients and their pets, without bias based on factors like owner income, ability/disability, animal breed or background. Clinic staff should retain oversight and be able to intervene when needed to uphold ethical standards in treatment decisions.

3 Transparency and explainability

From a veterinary perspective: If AI tools are being used in the clinic, all users, including clients, should be aware of the role that the AI is playing. The AI tool must be able to clearly explain how it arrived at its suggestions using plain language, in order for users to be able to question or challenge AI-based recommendations.

4 Robustness, security and safety

From a veterinary perspective: Veterinary AI systems must function reliably and accurately in the wide variety of conditions expected in the clinic (e.g., shortages in time or staffing, unexpected emergencies). They should be protected against cyber threats and include a way to safely override or be shut down if malfunctioning occurs. Regular risk assessments should be performed through the life of the AI tool to ensure ongoing safety for animals, clients and staff.

5 Accountability (for proper functioning of AI systems)

From a veterinary perspective: Clinic management and AI developers must clearly understand and communicate their responsibility for the AI's performance and ethical use. All decisions made with AI assistance should be traceable, enabling review if something goes wrong. A system should be in place that ensures users and providers claim responsibility for problems and manage risks. This should include cooperation between AI vendors, clinicians and support staff.

CHECKLIST

Download this checklist of important questions to ask when evaluating a new AI tool for your practice.



that is acceptable. With low-risk tools, you may be happy to communicate the presence of the risk to the staff using the tool and leave it at that. If you have limited time or budget, start with one or two of the highest-risk tools in your clinic and work from there.

Risk mitigation strategies will depend on the type of tool and the nature of the risk. Examples include:

- ▣ Increased communication with team members or clients
- ▣ Obtaining more in-depth informed consent from owners
- ▣ Having a human sign off on all decisions made by the tool
- ▣ Monitoring false positives or negatives and reporting them to the tool provider
- ▣ Limiting tool use to specific individuals or obtaining additional cybersecurity resources

Monitor Your AI and Your Governance Policy

AI technology is evolving incredibly fast. The possibilities and risks that it presents are evolving just as quickly. You will need to revisit your AI governance policy regularly, on a semi-annual to annual basis or whenever you encounter an AI technology that your policy doesn't cover. You'll also want to check on the performance of your existing stable of AI technology regularly. Over time, AI technology can experience model drift or staff may start using it for a purpose for which it wasn't intended.

Communicate with your team

Effective communication will make or break your AI governance program. AI technology will affect the entire team, so your entire team should be involved in the governance brainstorming from day one. As AI technology proliferates, members of your staff may learn about new tools before you do and bring great ideas to the table. You will be depending on your team to implement your responsible AI governance policy when it is completed. You will need their feedback on what is and isn't working, how their needs have changed or whether incidents have occurred.

Governance in a nutshell

We can, and must, implement AI technology in a responsible, trustworthy way today, not tomorrow. In doing so, we can help protect ourselves, patients, clients and the industry, while opening the door to an exciting new way of providing care.



Petra Harms, DVM

Dr. Harms is the founder and CEO of VetMaite, a veterinary AI consultancy and education platform. With over 15 years of clinical practice experience, her work focuses on responsible AI adoption, helping veterinary professionals build AI literacy and machine learning integration in food animal production. She holds certifications in AI and healthcare from Stanford and Harvard, and is trained in AI governance through the International Association of Privacy Professionals.

References

1. American Association for Veterinary State Boards. *Regulatory Considerations of the Use of Artificial Intelligence in Veterinary Medicine*. March 2025. Accessed April 2, 2025. <https://lp.constantcontactpages.com/sl/DGYUq08/AAVSBWhitepaperAIGuidance>
2. Artificial intelligence in veterinary medicine. Canadian Veterinary Medical Association. October 12, 2023. Accessed April 2, 2025. <https://www.canadianveterinarians.net/policy-and-outreach/position-statements/statements/artificial-intelligence-in-veterinary-medicine>
3. Appleby RB, Difazio M, Cassel N, Hennessey R, Basran PS. American College of Veterinary Radiology and European College of Veterinary Diagnostic Imaging position statement on artificial intelligence. *JAVMA*. 2025;263(6):773-776. doi:10.2460/javma.25.010027
4. Recommendation of the council on artificial intelligence. Organisation for Economic Co-operation and Development. Updated March 5, 2024. Accessed April 8, 2025. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
5. Acker J, et al. AIGP online training. International Association of Privacy Professionals. <https://iapp.org/train/privacy-training/OCT-AIGP>