Kicky van Leeuwen

Co-founder and consultant, Romion Health & Health AI Register, Utrecht, the Netherlands

> " I believe it is critical for healthcare organisations to understand the regulatory landscape when adopting AI "

Q1 What inspired you to transition from building AI models yourself to focusing on their validation and implementation in healthcare?

My background is in technical medicine, where I learned some programming but also spent several years in the clinic, shadowing doctors, much like medical students do. This combination was very insightful because, at the same time, I was developing AI models. However, I quickly realised the long and complex journey required to translate these models into clinically valuable tools.

One major challenge is regulatory clearance, which involves multiple steps before an AI model can even be used. Also, a model alone is not a complete solution or a software product. It requires many other steps to integrate it into hospital IT systems and make it truly functional in practice.

This realisation was a turning point for me. Understanding both the technical and clinical worlds, I wanted to bridge the gap between them. My focus shifted to ensuring that the AI research being conducted could actually reach clinical practice, benefit doctors, and, ultimately, improve patient care.

Q2 In your experience, what are the biggest barriers to the responsible adoption of AI-based medical software in healthcare, specifically radiology?

I think we've already made significant progress, which is a positive sign. In radiology, AI has been around for quite some time. When I started my journey, I worked in the radiology department at Radboud University Medical Center in the Netherlands because radiology was at the forefront of AI adoption. However, one of the biggest challenges remained the business case and financial sustainability of AI in healthcare.

AI should contribute to a more future-proof healthcare system that remains accessible to all who need it. However, for that to happen, it must also prove to be cost-effective in the long run. Right now, AI is often implemented as an 'add-on' to existing workflows, rather than replacing certain tasks. This raises the question: if AI increases costs, who is going to pay for it? In the future, AI should ideally shift towards an 'instead of' model, replacing certain tasks rather than just supplementing them. That would make the financial case clearer, but we're not quite there yet for many of the applications.

We also observe the financial challenges from the AI companies' side. Many AI vendors struggle to generate sustainable revenue, and we're seeing increasing market consolidation. Recently, one company ceased operations, and another was sold for only €1 million, far below its original investment. These financial struggles are deeply interconnected: if healthcare organisations find it difficult to adopt AI due to cost concerns, AI companies also struggle to sell their products at a sustainable price.

Q3 Could you describe how Romion Health, the company that you co-founded, helps healthcare organisations with AI strategy, procurement, and implementation?

With the rapid evolution of AI in healthcare, it can be challenging for caregivers and healthcare organisations to stay informed about available solutions and how to use them effectively. At Romion Health, our goal is to support organisations in the responsible adoption of AI in clinical practice. We help ensure they take an effective approach,

selecting solutions that truly fit their needs, rather than adopting AI simply because they want AI.

Responsible adoption also means ensuring that users have sufficient AI literacy and understand the associated risks. These risks can be financial, such as investing in AI without seeing the expected benefits, or clinical, where improper use could negatively impact patient care. Our mission is to minimise these risks as much as possible.

We do this in three main ways. First with consultancy and project management. We support healthcare organisations with AI implementation projects, helping them navigate the entire adoption process. Second, through education that increases AI Literacy. We provide education and training to healthcare professionals, ensuring they understand how to use AI safely and effectively. And third, through Health AI Register. This is a register of all the CE-marked AI solutions in healthcare to increase transparency in the field. This helps organisations gain a clear overview of available AI tools, understand their differences,

and make informed procurement decisions that align with their specific clinical needs.

Q4 What role do regulatory bodies play in ensuring that AI solutions are both safe and effective for clinical use?

During my PhD, I worked for a notified body, a regulatory organisation accredited by the European Union (EU) to grant CE marks for medical devices. AI software is often classified as a medical device, especially when it influences patient pathways or clinical decision-making. In these cases, regulatory bodies play a crucial role in determining whether AI solutions meet the necessary safety and performance standards.

For example, AI used in diagnosis, detection, or patient follow-up is considered a medical device. To enter the market, these solutions must obtain a CE mark, which requires manufacturers to provide extensive documentation to the notified body. This includes demonstrating how they ensure the quality of their AI solution, proving that its performance is on

par with the state of the art, and outlining measures to maintain its safety over time.

However, not all AI solutions in healthcare are classified as medical devices. A common example is AI scribes, which use voice-to-text technology to transcribe and summarise patient consultations. These systems are gaining popularity, but because they are considered administrative tools rather than medical devices, they do not fall under the same regulatory oversight. Unlike medical AI, AI scribes are not reviewed by a notified body (in the EU), an accredited body (in the UK), or the FDA (in the USA). This means there is no third-party authority ensuring their quality and reliability. As a result, healthcare organisations must take on the responsibility of assessing these AI tools themselves. They need to evaluate what mechanisms companies have in place to maintain the safety and effectiveness of their algorithms. I believe it is critical for healthcare organisations to understand the regulatory landscape when adopting AI solutions.

Q5 You have previously written about the potential of AI in breast imaging. Could you go into more detail about the current technologies that use AI for breast imaging, and what you think the future holds?

I find AI in breast imaging particularly exciting for two main reasons. First, breast cancer screening is a well-defined, single-task application. It is focused on one objective: analysing mammograms to detect potential malignancies and determining whether a woman should be referred to a hospital for further examination. This is a decision that radiologists traditionally make, but because it is such a specific and repetitive task, it is highly suited for AI. While we are now seeing AI models evolve to be more comprehensive, many of the current AI solutions remain highly task-specific, which makes this a promising area for AI adoption.

The second reason is the scale. Breast cancer screening programmes involve large populations, so implementing AI doesn't just benefit a small group, it impacts a vast number of people. In the Netherlands alone, around one million women are screened annually. This makes AI particularly valuable, as it can help optimise workflows and improve efficiency at a large scale.

Currently, there are multiple AI vendors with relatively mature solutions that are already being used in daily practice. Countries like Sweden and Denmark are leading the way in AI adoption for breast screening, implementing AI not just as an additional tool but in a way that actually reduces workload, an 'instead of' rather than an 'on top of' approach, as I mentioned before.

Traditionally, in Europe, two radiologists review each mammogram to determine whether further examination is needed. With AI, this process is changing. AI now serves as a second reader, meaning that in low-risk cases where AI and the first radiologist are confident that no cancer is present, a second human radiologist is no longer required. However, in higher-risk or uncertain cases, the traditional two-radiologist approach is maintained. This shift is particularly exciting because it represents a true efficiency gain, moving towards 'not doing something anymore' rather than simply adding AI on top of existing workflows.

We are thrilled that with our company, we lead the implementation of AI in a national breast cancer screening programme, and we can work on such a high-impact application of AI in healthcare.

Q6 What are some common misconceptions about AI in radiology that you frequently encounter?

One of the biggest misconceptions I still encounter is the belief that using AI in radiology means that all patient data will be shared with the vendor and used to retrain their algorithm. In reality, this is not the case unless specific agreements (such as collaboration agreements) explicitly allow data sharing. By default, patient data remains with the hospital. In some cases, data may need to be temporarily shared for processing, such as when using a cloud-based solution, but that does not grant vendors the right to use it for further development or retraining. This is an important distinction that not everybody is aware of.

Another common misconception is that AI models are continuously retrained in real-time based on user interactions and feedback. In practice, this is quite limited. AI algorithms are typically deployed in specific versions that are periodically updated. A hospital may receive a new version of the algorithm, but in most cases, individual hospitals do not have their own AI models that are continuously retrained while in use. There are exceptions though. For example, in breast cancer screening, one company allows authorised users to provide feedback that can be incorporated into model retraining. However, this is the exception rather than the norm.

A regulatory misconception, particularly common outside of radiology, is the assumption that having a human in the loop exempts an AI solution from medical device regulations or the AI Act. Some vendors claim that because a human is required to approve AI-generated decisions (e.g., by clicking a button), their solution does not qualify as a medical device and does not require regulatory oversight. However, this is not how medical device regulations or AI-related risk assessments work. While human oversight is an important safeguard, it does not eliminate potential risks or remove the need for regulatory scrutiny.

Q7 What developments in AI for radiology are you most excited about in the coming years?

One of the most exciting developments in AI for radiology is the growing focus on post-market surveillance and post-deployment monitoring. Post-market surveillance is already a regulatory requirement for AI vendors: they must continuously monitor their solutions once they are deployed in clinical practice. However, healthcare organisations also share the responsibility of assessing how these AI solutions perform in their specific settings, which is why the term 'post-deployment monitoring' is increasingly used.

This is a crucial prerequisite for moving toward more autonomous AI. If we want AI to take on greater responsibility in decision-making, potentially without a human involved in every single decision, there must still be oversight to ensure it operates safely and as expected. Monitoring is essential to detect deviations caused by factors such as changes in data distribution, modifications in imaging acquisition (e.g., the introduction of a new machine), shifts in patient populations, or AI system updates that may unexpectedly alter performance.

To support this, AI platforms and AI marketplaces are beginning to offer tools, dashboards, and solutions for more effective monitoring. As these technologies continue to mature, I look forward to seeing an increasing responsibility for AI, while ensuring that AI-driven decision-making remains safe, transparent, and accountable.

Q8 If you could give one piece of advice to healthcare organisations looking to adopt AI, what would it be?

AI is here to stay, and ideally, it will become so seamlessly integrated into healthcare systems and workflows that it eventually feels like just another tool, something you don't even consciously notice. As AI becomes more embedded, it remains essential for healthcare organisations to build knowledge and expertise in this area.

That doesn't mean every healthcare professional needs to know how to build or program AI models, but they do need to understand how to use AI responsibly and be aware of its risks. I see a strong parallel with cybersecurity. In an organisation, not everyone is a cybersecurity expert, but staff members are still expected to follow basic security practices. Many organisations provide annual e-learning modules or phishing awareness exercises to ensure employees understand their role in maintaining security.

I believe we need a similar approach for AI: ongoing education to keep healthcare professionals up to speed on their responsibilities. This includes understanding potential biases in AI outcomes and recognising when AI suggestions should be questioned. By ensuring sufficient AI literacy across healthcare teams, we can ensure that AI is used responsibly, ultimately improving patient safety and clinical decision-making.