



> Artificial Vigilance

Where AI Meets Drug Safety



essjay
SOLUTIONS LTD

> What can UK AI Healthcare Regulation Learn from The Pharmacovigilance Approach?

Pharmacovigilance has long operated on the principle that regulatory approval is not the end of safety assessment, but the beginning of **continuous monitoring**. Now, the UK Government is moving towards applying a similar philosophy to **AI in healthcare**.

The UK Government has announced it will **accept** all **44 recommendations** made by the **National Commission** into the **Regulation of AI in Healthcare**, signalling a shift in how AI may be regulated across the NHS. Their response outlines ambitions to replace traditional one-off assessments with a **flexible framework** that continuously monitors AI systems throughout their lifecycle.

The independent NHS doctor-led commission developed its recommendations through consultation with over 12,000 **patients, clinicians** and members of the **public**. It argues that existing regulatory approaches must **evolve** to address fluctuating AI performance after deployment, making static assessments inadequate.

The proposed lifecycle-based framework would emphasise evaluating AI systems in **real-world clinical environments**, rather than relying predominantly on pre-deployment assessments. The MHRA is already exploring this approach through its AI Airlock regulatory sandbox, focusing on **post-market surveillance** of AI-enabled medical devices.

These recommendations recognise that AI systems are **dynamic**. Performance can **change** as models are updated, clinical contexts evolve and new data becomes available. Consequently, demonstrating safety before deployment does not guarantee continued reliability. For PV professionals, this should sound familiar.

Medicines undergo rigorous pre-market evaluation, but monitoring does not stop there. Safety profiles continue to be assessed through real-world evidence, adverse event reporting and ongoing benefit-risk analysis. Regulators increasingly recognise that AI may require comparable **post-deployment surveillance**.

“I particularly welcome the commitment to monitoring AI throughout its use in practice, strengthening patient involvement in decision-making and ensuring there are clear routes to redress when things go wrong”

Henrietta Hughes, Patient Safety Commissioner for England

The third phase of the MHRA's AI Airlock will investigate **lifecycle regulation** of AI-enabled medical devices to inform future guidance. Further commitments include guidance on managing **changes** to AI devices and exploring **staged authorisation** pathways permitting supervised clinical deployment while additional evidence is gathered.

These PV parallels invite an intriguing possibility: rather than AI simply transforming pharmacovigilance, could pharmacovigilance's established **safety principles** help **shape AI regulation**? Concepts such as signal detection, continuous benefit-risk assessment and systematic safety reporting could offer valuable lessons for **identifying emerging problems** in AI systems.

For an industry accustomed to asking whether AI can improve pharmaceutical safety, it is worth asking whether the principles of pharmaceutical safety can improve AI.

> Could AI Predictions Make Polypharmacy Pharmacovigilance Prophylactic?

What if pharmacovigilance could identify **dangerous drug interactions** before patients ever experience them? Recent breakthroughs suggest AI could predict and potentially prevent adverse reactions from drug combinations.

DCSE (**Drug Combinations Side Effects**) is a machine learning model developed by researchers Ruben Jimenez and Alberto Paccanaro to predict adverse reactions from combinations of different medicines. They seek to address a longstanding challenge in drug safety: clinical trials realistically **cannot investigate every** possible drug **combination**, so many potentially harmful interactions only emerge when these medicines are combined in larger populations.

“DCSE-twin represents an opportunity to use patient population data to flag and prioritise adverse effects caused by drug combinations for further validation”

Dr Ruben Jimenez, researcher and co-author of the 2026 DCSE study

DCSE works by learning numerical representations, called latent signatures, of individual medicines, drug combinations and adverse events. By identifying known patterns within existing drug safety data, the model can **estimate** the **likelihood** of **specific side effects** occurring when two medicines are taken together. DCSE can even make predictions for drug combinations with no previously documented adverse reactions.

To test its predictive capabilities, researchers trained the model on historical data and assessed whether it could predict side effects reported between 2009 and 2014. Impressively, DCSE **consistently outperformed** competing methods, even when evaluating previously uncharacterised combinations.

Although traditional pharmacovigilance has focused on identifying safety signals from adverse reactions that have already happened, predictive models like DCSE could invite an additional **preventative layer**, allowing safety teams to investigate potentially dangerous combinations before patient exposure.

These predictions could inform **risk mitigation** strategies, **guide the design** of clinical trials and **prioritise combinations** requiring further investigation. Instead of being faced with the insurmountable task of investigating every possible combination, clinical trials could just investigate combinations which algorithms predict could be problematic. Perhaps in the future, validated systems could also **support prescribing** decisions and **warn clinicians** about potential interactions before treatment begins.

Of course, these predictions are not gospel. A model's identification of an association cannot establish causality, and reliable clinical deployment requires much more validation and careful consideration of false alarms. However, perhaps the most interesting insight from this innovation is the possibility of pharmacovigilance evolving from a mainly **reactive** discipline into an **anticipatory** one. The ultimate ambition is not simply to detect adverse reactions faster, but to stop them from happening at all.

> Right Patient, Right Trial: How AI Can Transform Clinical Trial Recruitment

Finding the best patients for clinical trials is one of the pharmaceutical industry's most persistent challenges. **Recruitment** can take months or even years, **delaying research** and limiting patients' access to investigational treatments. However, AI may hold the answer to make the process **faster**, more **efficient** and potentially more **inclusive**.

“Anyone who's considering a clinical trial is making a big decision. AI can help us create more awareness, identify eligible patients, and streamline the process”

Melinda Rottas, Head of Optimisation, Analytics and Recruitment Solutions at Pfizer

Pfizer is experimenting with **AI-assisted recruitment** to identify eligible patients, optimise the locations of the trials and streamline the design of studies. Rather than relying on manual review of medical records, AI can analyse clinical information against **complex eligibility criteria** and identify candidates who might otherwise have been overlooked.

The potential benefits are supported by independent research. Researchers from the US National Institutes of Health developed TrialGPT in 2024, an AI system that matches patient information with clinical trial criteria. TrialGPT achieved **87.3% accuracy** when assessing individual eligibility criteria and in a small pilot, clinicians reported a **reduction in screening time** of **42.6%** without a significant difference in accuracy.

For pharma companies, **greater recruitment efficiency** could directly lead to **faster trial completion**, **reduced administrative burdens** and **shorter drug development timelines**. Furthermore, AI could help address **inequalities** in clinical research. Traditional recruitment often relies on established research centres and referral networks. By analysing where eligible patient populations receive care, AI could help sponsors identify overlooked communities and select more **representative** trial sites.

More diverse clinical trial populations generate safety evidence across a broader range of **demographic** and **clinical characteristics** and will improve understanding of how medicines affect different groups. This is not to say AI is not susceptible to **bias**; indeed, systems trained on incomplete or unrepresentative healthcare data could reinforce existing inequalities rather than mitigate them. But if implemented responsibly, AI could help clinical research become **faster**, more **accessible** and more **representative**.

For pharmacovigilance professionals, that last benefit may prove particularly valuable: better representation during development could mean better safety knowledge when medicines reach the wider population, yielding drug safety **benefits** beyond simple efficiency and automation.



> When the AI Bubble Bursts, What Happens to Pharmacovigilance?

The AI boom has attracted **trillions** in investments, but growing concerns about **debt**, **profitability** and **unsustainable spending** are raising an uncomfortable question: what happens if the **AI bubble** bursts? For PV professionals increasingly reliant on AI-powered systems, the consequences could extend beyond financial markets, and many market analysts seem to see the AI bubble as more of a question of when rather than if it will burst.

Technology giant **Oracle** serves as a concerning example. According to Reuters, the company had accumulated **\$129.5 billion** in **debt** by June 2026, alongside **\$260 billion** in long-term data centre **lease commitments**. Its aggressive leveraged investment strategy has led to its credit being downgraded to **BBB-**, just one level above junk status.

Oracle is also invested in drug safety. Its **Argus safety** platform supports adverse event processing, safety data management and regulatory reporting. They recently announced a **partnership** with **Wipro** to strengthen AI-powered PV services using its cloud-based Argus business services model. Clearly there is an increasingly **interconnected** relationship between PV operations and the technology companies whose financial fortunes are tied to the wider AI economy.

The concerns extend beyond Oracle. Broadcom CEO Hock Tan, when questioned on whether companies were generating meaningful returns from their AI investments, warned that **“most companies out there cannot find any return on investment”**.

So what happens if the bubble bursts? While a collapse in AI valuations would not instantly disrupt PV, a prolonged **investment downturn** could have more serious consequences. Reduced funding could force smaller AI vendors to **scale back** development, **increase prices** or **cease operations**. PV companies relying on these vendors may have their automation projects **disrupted**, **lack technical support** or face **unexpected costs** to replace systems. More seriously, financial distress affecting providers of safety infrastructure could create **business continuity**, **data accessibility** and regulatory **compliance challenges**.

The irony is that AI promises to help pharmacovigilance anticipate and manage risks, yet the AI industry's reliance on debt and speculation may have introduced a brand-new risk that PV companies must carefully monitor. AI isn't going anywhere; the dotcom bubble didn't delete the internet. But perhaps it is now worth considering not just the practical benefits of an AI system, but the financial sustainability of its provider.

“The AI boom is a bubble of epic proportions, and there's no way to stop it from popping and taking down the stock market and economy with it”

Michael Burry, “The Big Short” Investor



essjay
SOLUTIONS LTD

Contact Us

For further information:

 www.essjaysolutions.co.uk

 sabita.mukherjee@essjaysolutions.co.uk
