KPMG

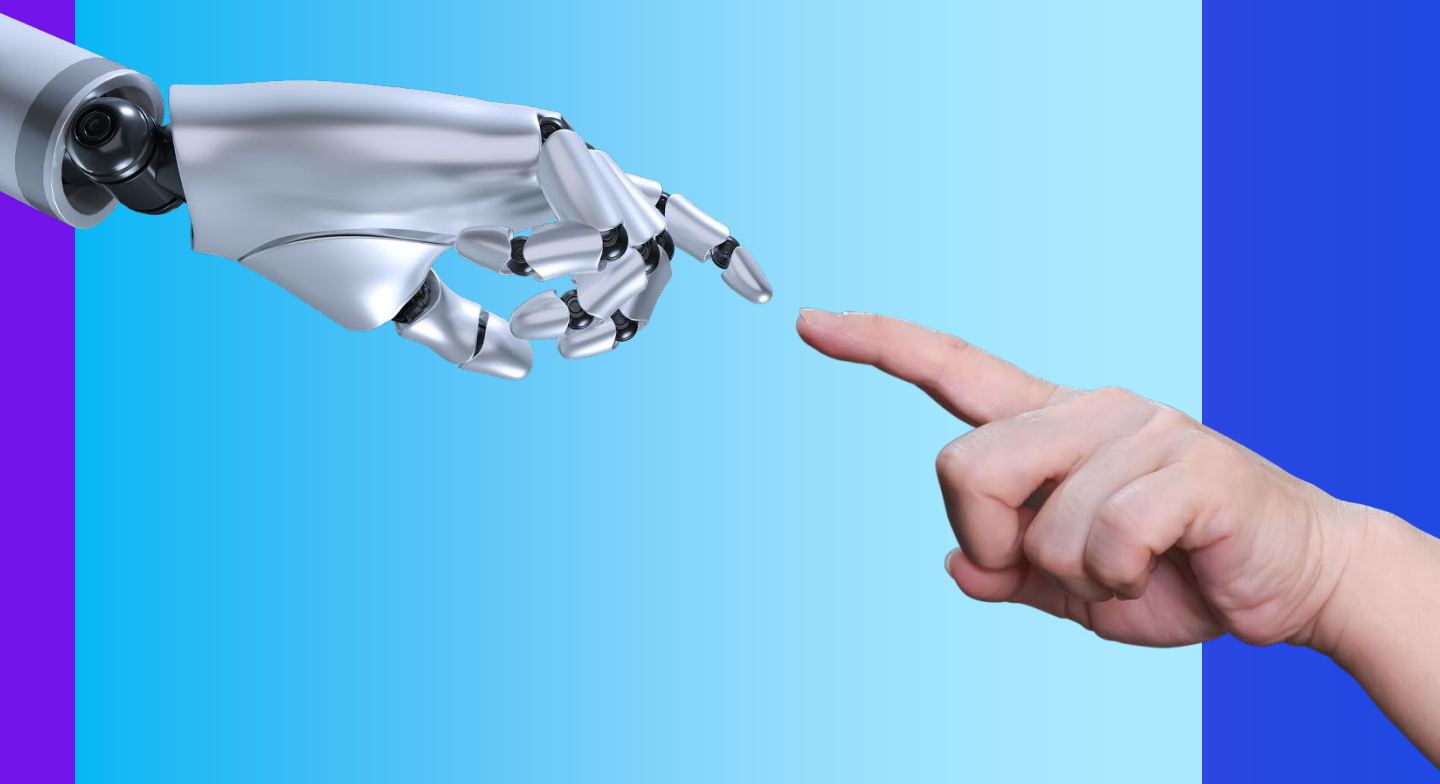

Building trust in artificial intelligence: the journey to effective regulation

KPMG Life Science regulatory solutions: Ewan Kerr-Edwards

—

July 2023

The risk/benefit paradigm of AI

Developments in Artificial Intelligence (AI) have accelerated at an unprecedented pace with innovation progressing exponentially. As society considers the impact of AI on jobs, productivity and decision-making, a dichotomy of thought begins to emerge.

On the one hand, AI is viewed as the panacea to many of our problems from workforce shortages & healthcare delivery to productivity challenges. On the other hand, AI can be viewed as an existential threat that must be understood, mitigated and regulated. This article will explore the risk benefit paradigm presented by AI and how stakeholders may navigate this brave new world and why the future of AI adoption will involve building societal trust and appropriate controls.

Advances in Generative AI have accelerated the adoption of AI technologies in all industries with businesses seeking to implement Large Language Models (LLMs) into their processes to improve efficiency. Over the next decade, companies will begin to adopt LLMs en masse to unlock productivity and reduce the time spent on repetitive tasks. However, this scale of adoption presents emergent risks over the veracity of AI predictions and integral bias that can plague machine learning algorithms. Regulators will be keen to understand this technology better and adapt their systems to ensure that AI is used responsibly and without bias. The UK has signaled that no new regulator[1] will be created to oversee AI and no legislation is currently planned to regulate the use of AI meaning existing regulators will be required to monitor these systems, potentially leading to a crossover of multiple jurisdictions and a lack of clear accountability for the control of rapidly evolving AI technology. The NHS transformation Directorate have recently launched a multi-agency digital regulations service[2] that is an good example of cross-agency collaboration that is taking place in the regulatory space.

Note: [1] https://www.reuters.com/world/uk/britain opts adaptable airules with no single regulator 2023 03 28/

[2] https://www.digitalregulations.innovation.nhs.uk/

What will good AI regulation look like?

The step change improvements in AI has ushered in calls for regulation. A question that often arises is whether we need to regulate such a nascent technology which has been hotly debated by experts in recent months. What they all agree on is the enormous impact AI is likely to have on society and that this will only accelerate over the coming decades. Regulation of AI is an important step in building trust in new technologies and enabling them to be implemented into important societal decision making whilst demonstrating their safe and effective use.

In order to begin talking about regulation of AI, we must first deconstruct what good regulation looks like. Good regulation such as those for Healthcare products, Aircraft and Nuclear technologies, approach the activity in a similar way. Providing a clear and transparent framework that is proportional and risk based ensures public protection and provides room for innovation in the private sector. International harmonization will be key to standardising the approach to AI regulation in the future.

When regulating AI, there are 3 key focus areas that can provide a preliminary framework:

1. **Risk-based approach to the classification of AI.** Stratifying use cases by risk can lead to a risk-based approach to compliance and approval. AI that is used to affect patient treatment decisions should be treated with the necessary controls to ensure safety.

2. **A transparent framework for compliance and enforcement** – regulators may need to commission AI experts to examine tools that perform fundamental societal tasks to truly understand the way AI makes decisions

3. **Versatility of regulation to promote innovation –** we have seen from the pandemic how regulators can adapt to a rapid influx of data to approve vaccines. A similar speed and brevity will be needed in regulatory approval of AI products due to the fast-paced nature at which these innovations are happening.

Global AI regulatory approach

The U.S. National Institute of Standards and Technology (NIST), a federal agency of the Department of Commerce released it's voluntary Risk Management Framework for AI in January 2023[3] to align technical development of risk management frameworks across industries. The process builds on the OECD's Framework for Classification of AI[4] systems that puts People & Planet at the centre. These help to identify the risks and classifications of AI systems throughout the AI product lifecycle whilst stopping short of the EU AI Act's legislative implementation. Providing a framework for AI Risk Management that takes into account impacted stakeholders, the economic context and the AI model deployment will assist in the coordination of international standards for AI.

The EU AI act sets out an initial framework for how AI might be regulated in practice. The draft of the regulation seeks to determine the risk of AI applications including a section deemed unacceptable (Title II) such as social scoring and remote biometric identification in publicly available spaces. Title III, Chapter II, Article 9 stipulates the risk management system, data governance, human oversight and record keeping requirements for High Risk AI applications (e.g. medical devices). This set of regulatory requirements are a good indicator of how AI companies will be impacted by regulation in the coming years. Stakeholder engagement is crucial for AI companies beginning to think about compliance early as this will have a positive impact on future regulatory alignment and business growth.

OECD AI classification system[4]

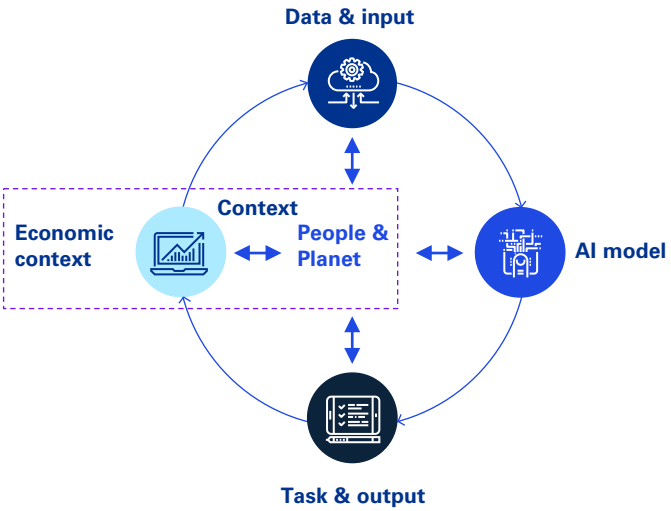

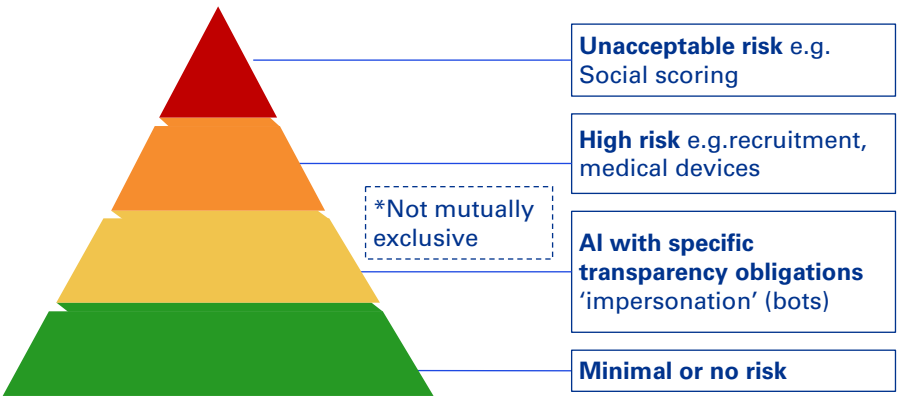

Unacceptable risk e.g. Social scoring	**Prohibited**
High risk e.g.recruitment, medical devices	**Permitted** subject to compliance with AI requirements and ex-ante conformity assessment
AI with specific transparency obligations 'impersonation' (bots)	**Permitted** but subject to information/transparency obligations
Minimal or no risk	**Permitted** with no restrictions

Note: [3] NIST - Artificial Intelligence Risk Management Framework (AI RMF 1.0)
[4] SOECD Digital Economy Papers No. 323 OECD FRAMEWORK FOR THE CLASSIFICATION OF AI YSTEMS
[5] European Commission guidance document

High risk AI & the EU AI act

The European Union is developing the AI Act[6] to regulate High-Risk Systems. The European Parliament voted through the draft text in March 2023 with expected signature into the Official Journal in 2024 or 2025. Below describes the obligations of businesses developing high-risk AI:

High-risk AI systems

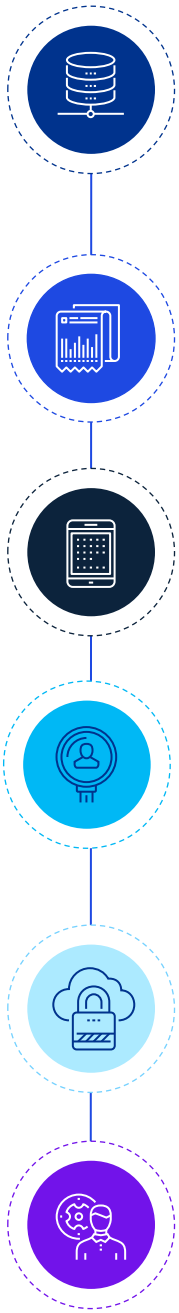

Data and data governance (article 10)

Training, validation and testing data sets shall use appropriate data governance practices:

- **Design choices**, **data collection**, **data prep-processing** e.g. annotation labelling and aggregation
- Suitability, availability and quantity of data sets
- Evaluation of possible **biases** & specific geographic, behavioural and functional settings

Technical documentation (article 11) & record keeping (article 12)

Technical Documentation will be drawn up **prior** to high-risk system being placed on the market (Annex IV) including:

- General Description of system, **Detailed Description** of AI System and **Development** Process
- Details of changes made to system, harmonised standards applied & **EU Declaration of Conformity**

Automatic recording of events (logs) and **traceability** of functioning

Transparency and provision of information to users (article 13)

Operation shall be **transparent** and enable users to interpret the output and use it appropriately

- Shall include instructions for use that include capabilities, limitations
- Lifetime of the AI system, maintenance and **software updates**

Human oversight (article 14)

Shall be developed so that there are appropriate human interface and can be overseen by natural persons:

- Will be aimed at **reducing risks** to health, safety and human rights and put in place prior to marketing
- Humans given oversight to fully understand **capacities** and **limitations** of AI system, remain aware of over-reliance (automation bias)

Accuracy, robustness and cybersecurity (Article 15)

High risk AI systems will maintain, accuracy, robustness and cybersecurity through the lifecycle:

- Level of accuracy stated in the **instructions for use**
- Resilient to errors with back up plans or fail-safe plans in place
- Must be resilient to **unauthorised 3rd parties** attempt to access the system in line with current standards

Risk management system (Article 9)

Continuous and iterative process through entire lifecycle

- Identification and Risk Analysis of all **known and foreseeable** risks
- Estimation of risks under reasonable foreseeable misuse vs the **intended purpose**
- Evaluation of emergent risks identified by Post-Market monitoring

All organisations will have to be aware of their obligations with regards to data that they have collected for use in machine learning algorithms. In the Life Science Sector, Data Governance, Transparency and Risk Management are established principles that will have to evolve to address the convergence of healthcare regulation with AI such as the EU AI Act. Data accumulated from research or real world use of assets; E.g. large repositories of biological data that can be used by Machine Learning algorithms to derive meaningful insights will require a robust data governance strategy.

Note: [6] EU AI Act (2021) Regulation Of The European Parliament And Of The Council Harmonised Rules On Artificial Intelligence (Artificial Intelligence Act)

AI safety research

AI Alignment is an area of AI safety research that aims to ensure AI systems meet their desired outcomes based on ethical principles and human-stated goals. LLMs have the capacity to generate untruthful, toxic and harmful sentiments[6] which becomes especially prescient in a clinical context where answers generated have a direct patient impact. Privacy and confidentiality of patients' clinical and medical data remains a significant area of concern, and strong security controls should be in place for the training models to retain their integrity.

In their current state, LLMs capability in clinical decision making is unknown, given the risks posed by incorrect answers generated in response to questions, the potential for unauthorised changes to the data in the training models and lack of representative clinical data. However, there has been progress; Google's Med-PaLM 2 recently achieved 85.4% on the U.S. Medical Licensing examination[7] signaling strong progress in medical applications. Meanwhile, ChatGPT, which is underpinned by OpenAI's GPT3.5 was trained using reinforcement learning from human feedback (RLHF) has shown initial signs of supporting scientific research and extracting useful insights from Electronic Health Records (EHRs)[8].

The MHRA have stated that the intended use of the LLM in a clinical sense will determine the applicability of the Medical Device Regulations in the UK[9]. Defining the specific intended use of an LLM can be tricky due to the broad and versatile nature of the model, however the UK government has also recently released guidance[10] on how to specifically craft an intended use statement for software products intended to be used in a medical context to make this process easier. In this way regulation can be used as an enabler for innovative businesses who require extra guidance as AI is increasingly used in medical applications. Ensuring that the intended patient population, clinical task performed and operating environment are well defined is fundamental to establishing AI products as medical devices in the UK. Software used as a medical device traditionally follow the IEC 62304 series of standards for compliance which has proved an effective enabler of documented evidence of compliant software engineering practices. In future, the use of standards such as these will need to be flexible to account for the rapid development of AI.

Note:

[6] https://openai.com/research/instruction-following

[7] https://sites.research.google/med-palm/

[8] Cascella M, et al. Evaluating the Feasibility of ChatGPT in Healthcare: An Analysis of Multiple Clinical and Research Scenarios. J Med Syst. 2023 Mar 4;47(1):33. doi: 10.1007/s10916-023-01925-4.

[9] https://medregs.blog.gov.uk/2023/03/03/large-language-models-and-software-as-a-medical-device/

[10] https://www.gov.uk/government/publications/crafting-an-intended-purpose-in-the-context-of-software-as-a-medical-device-sand/crafting-an-intended-purpose-in-the-context-of-software-as-a-medical-device-samd

AI alignment and use cases

AI Life Science Use cases

In the last 10+ years there has been an acceleration in novel genomic sequencing techniques that has seen a rapid uptick in levels of data generated in cell biology. With greater tools for data capture, the industry now has access to burgeoning levels of data relating to the way cells behave, grow and express proteins. Tools such as RNA-seq and Next Generation Sequencing provide data on genomics, epigenetics and transcriptomics that will power forward the use of biomarkers in both medicines and diagnostics. AI is starting to be used as a tool to derive relationships between gene expression levels and disease whilst concurrently the regulations of diagnostics have been overhauled in the EU. Companies needing to navigate the new system of classification, clinical performance evaluation and post-market surveillance requirements to meet the new standards are looking to the U.S. to launch new devices where the regulatory system is more stable. The UK is well placed to capitalise with innovative regulatory science applied to AI in *in vitro* diagnostics.

AI company spotlight - Ultromics

Ultromics is a pioneering company in the AI space that has made significant strides in the field of medical imaging, a spin-out from Oxford University, they utilize AI to diagnose Heart Disease in Ultrasound scans of the heart. Ultromics has received four FDA approvals and two CE marks and recently also received two breakthrough device designations from the FDA for its AI-powered diagnostic tools.

Ultromics has managed to clear a new product each year through the FDA, this impressive pace is largely due to investing in building a solid relationship with the FDA but also providing rigorous assessment of the product efficacy and using large, generalized data-sets for validation

Looking ahead, a major challenge in the EU is regulator capacity and availability of notified bodies to support audits. Also the new MDR requirements for post-market clinical monitoring are quite onerous for start-up companies. There are however encouraging signs that the FDA are starting to better understand the AI ecosystem and invest in clearing new algorithms, they see the interaction with companies like Ultromics as a good way to learn and improve their capabilities but they still struggle with assessing the risk of using AI in clinical practice. Ultromics approaches this by focusing on the relative risk reduction of using computer-aided diagnostics versus human-centred analysis which is error prone and operator dependant but also reinforcing that the AI is an adjunct diagnostic support rather than replacing the decision making of clinicians. The future of AI regulation is still uncertain, but it is clear that there is a need for more collaboration between regulators and industry players to build the right regulatory framework that balances speed and impact with risk.

\- Ultromics COO Amir Hasan

Conclusion

Artificial Intelligence is developing rapidly. AI will inevitably reach all aspects of society including healthcare with the public rightly wanting more clarity on how these systems will be used and regulated. With recent advancements in science and technology, the opportunity to apply AI to healthcare problems has never been greater. However, risks remain around the clinical efficacy of this technology, how value is captured from the healthcare data we produce and maintaining the public's trust in the models and their output.

Regulation in this field is nascent, however as the power of AI is increasingly demonstrated, and societally important decisions are taken by these systems then regulation will be required to ensure safety and effectiveness of this technology. Transparency of decision making and training is the key to establishing a robust and streamlined regulatory framework. Appropriate controls are required to protect individuals' privacy and protect the integrity of the data so the public can maintain trust in the model output. Regulators will need to develop a deep understanding of AI and Machine Learning models. Pragmatic and effective regulation is critical to establish a safe and effective AI ecosystem and is crucial to build societal trust and adoption of AI.

KPMG Life Science Advisory Teams are helping companies to navigate the emerging AI regulatory landscape. We support companies to transform their organisations ensuring data governance, regulatory compliance, transparency, cybersecurity and risk management are embedded across all functions.

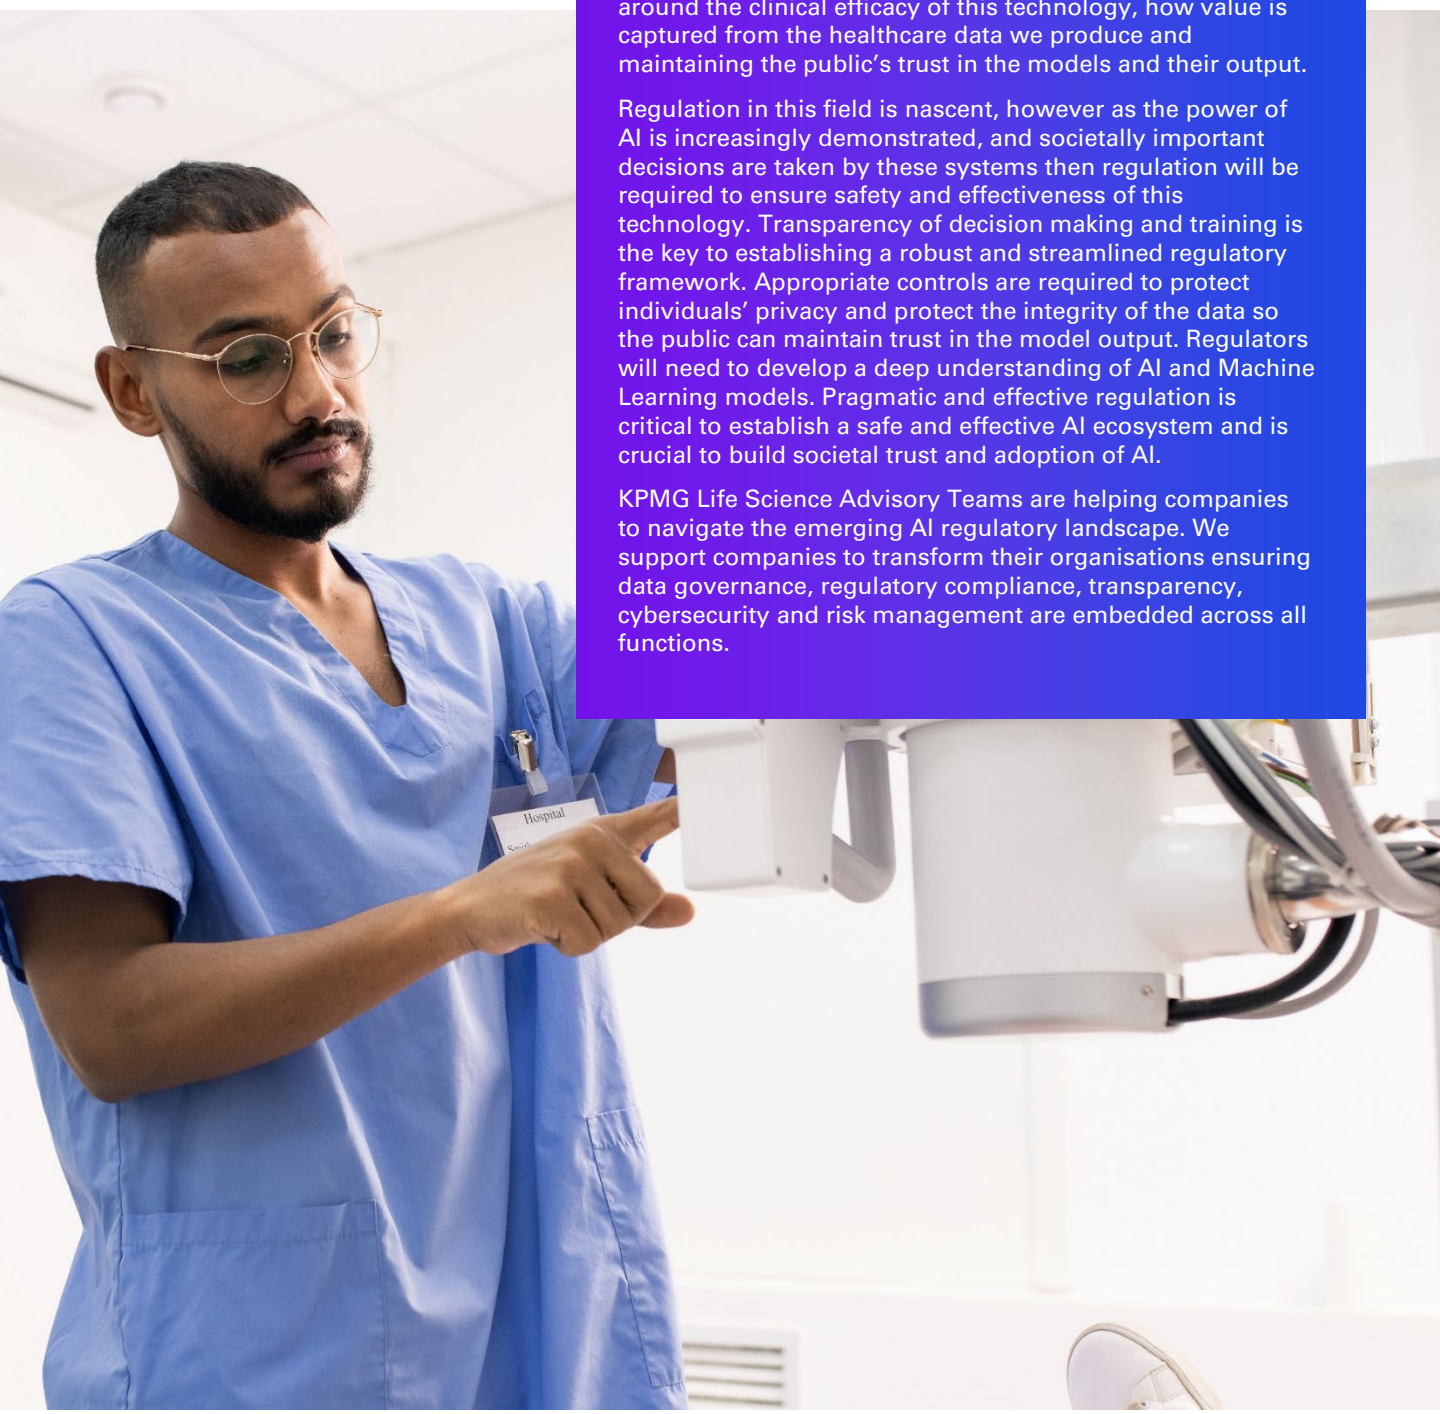

Contact:

Adrian Griffiths
Healthcare and Life Science Lead
KPMG in the UK
T: +44 (0)771 727 2072
E: adrian.griffiths@kpmg.co.uk

Anusha Foy
Life Science and Biotech Regulatory Solutions Lead
KPMG in the UK
T: +44 (0) 7510 376178
E: anusha.foy@kpmg.co.uk

Caroline Rivett
Cybersecurity & Privacy Lead
KPMG in the UK
T: +44 (0) 7990 577427
E: caroline.rivett@kpmg.co.uk

Varun Sukumaran
Senior Manager
KPMG in the UK
T: +44 (0) 7543509754
E: varun.sukumaran@kpmg.co.uk

Phil Brame
Senior Manager
KPMG in the UK
T: +44 (0) 7935 603347
E: philip.brame@kpmg.co.uk

Ewan Kerr Edwards
Associate Manager
KPMG in the UK
T: +44 (0) 7510 376758
E: ewan.kerredwards@kpmg.co.uk

kpmg.com/uk

Some or all of the services described herein may not be permissible for KPMG audited entities and their affiliates or related entities.

Document Classification: KPMG Public

CREATE | CRT150245B