

Bristows

Spotlight on...

The EU AI Act: what you need
to know right now

Introduction

The EU AI Act is a game-changing AI regulatory framework that will be significant for organisations inside and outside the AI industry. By now, many organisations will be implementing AI governance frameworks to help them comply with the AI Act, and many lawyers will have received training on the Act and many are leading their organisation's compliance efforts.

At Bristows we have been advising clients on a range of issues relating to the AI Act: how the Act classifies their AI products and services, helping design governance frameworks mapping the roles and responsibilities needed, and how they can work and contract with third party AI providers, all with effective risk management in mind.

A common question we get is, "what do I need to know about the AI Act *right now*?" Essentially, what are the core components of the Act, how might it affect my organisation, and what are some useful ways to better understand my role in advising on AI?


So, in this article series, our experts unpack what we think organisations across all sectors should understand about this significant regulatory change to inform their compliance strategy.

Contents




[Is my AI "high-risk" under the AI Act?](#)

Much of the Act is focused on defining and regulating "high-risk" AI systems and use cases, so this is crucial to understand given the requirements that flow from this question.



[What is my role under the AI Act?](#)

The Act does not only regulate "Big Tech", all manner of organisations developing, building upon and using AI are potentially caught.



[Are data protection governance frameworks a good model for AI governance?](#)

Lessons from a recent significant regulatory change: GDPR and data protection.



[The EU AI Act as product safety legislation – lessons from the world of medical devices](#)

This significant aspect of the AI Act has close parallels with existing product safety laws, so we look at lessons from medical device regulation.



[Copyright issues under the EU AI Act](#)

IP is one of the main issues for everyone involved in the burgeoning AI ecosystem, so we look at how the Act deals with copyright issues.



[What employers need to know about the EU AI Act](#)

As many internal AI use cases will affect employees, we look at what employers need to know about the Act.

Meet the team



Vik Khurana
Partner, IT & digital
[Visit profile](#) | [Email](#)



Alex Denoon
Partner, Regulatory
[Visit profile](#) | [Email](#)



Marc Dautlich
Partner, Data Protection
[Visit profile](#) | [Email](#)



Toby Headdon
Of Counsel, Brands, Designs & Copyright
[Visit profile](#) | [Email](#)



Charlie Hawes
Of Counsel, IT & digital
[Visit profile](#) | [Email](#)



Emily Atkinson
Associate, Employment
[Visit profile](#) | [Email](#)

Is my AI “high-risk” under the AI Act?

The primary purpose of the AI Act is to prevent risks caused by “high-risk” AI systems. Yes, the Act does other things. It bans some AI use cases. It imposes transparency measures on deceptively realistic AI. It has a standalone section for foundation models, with additional safety rules designed for future generations of increasingly powerful models. But most of the Act is a legislative framework for the regulation of AI systems that it classifies as high-risk.

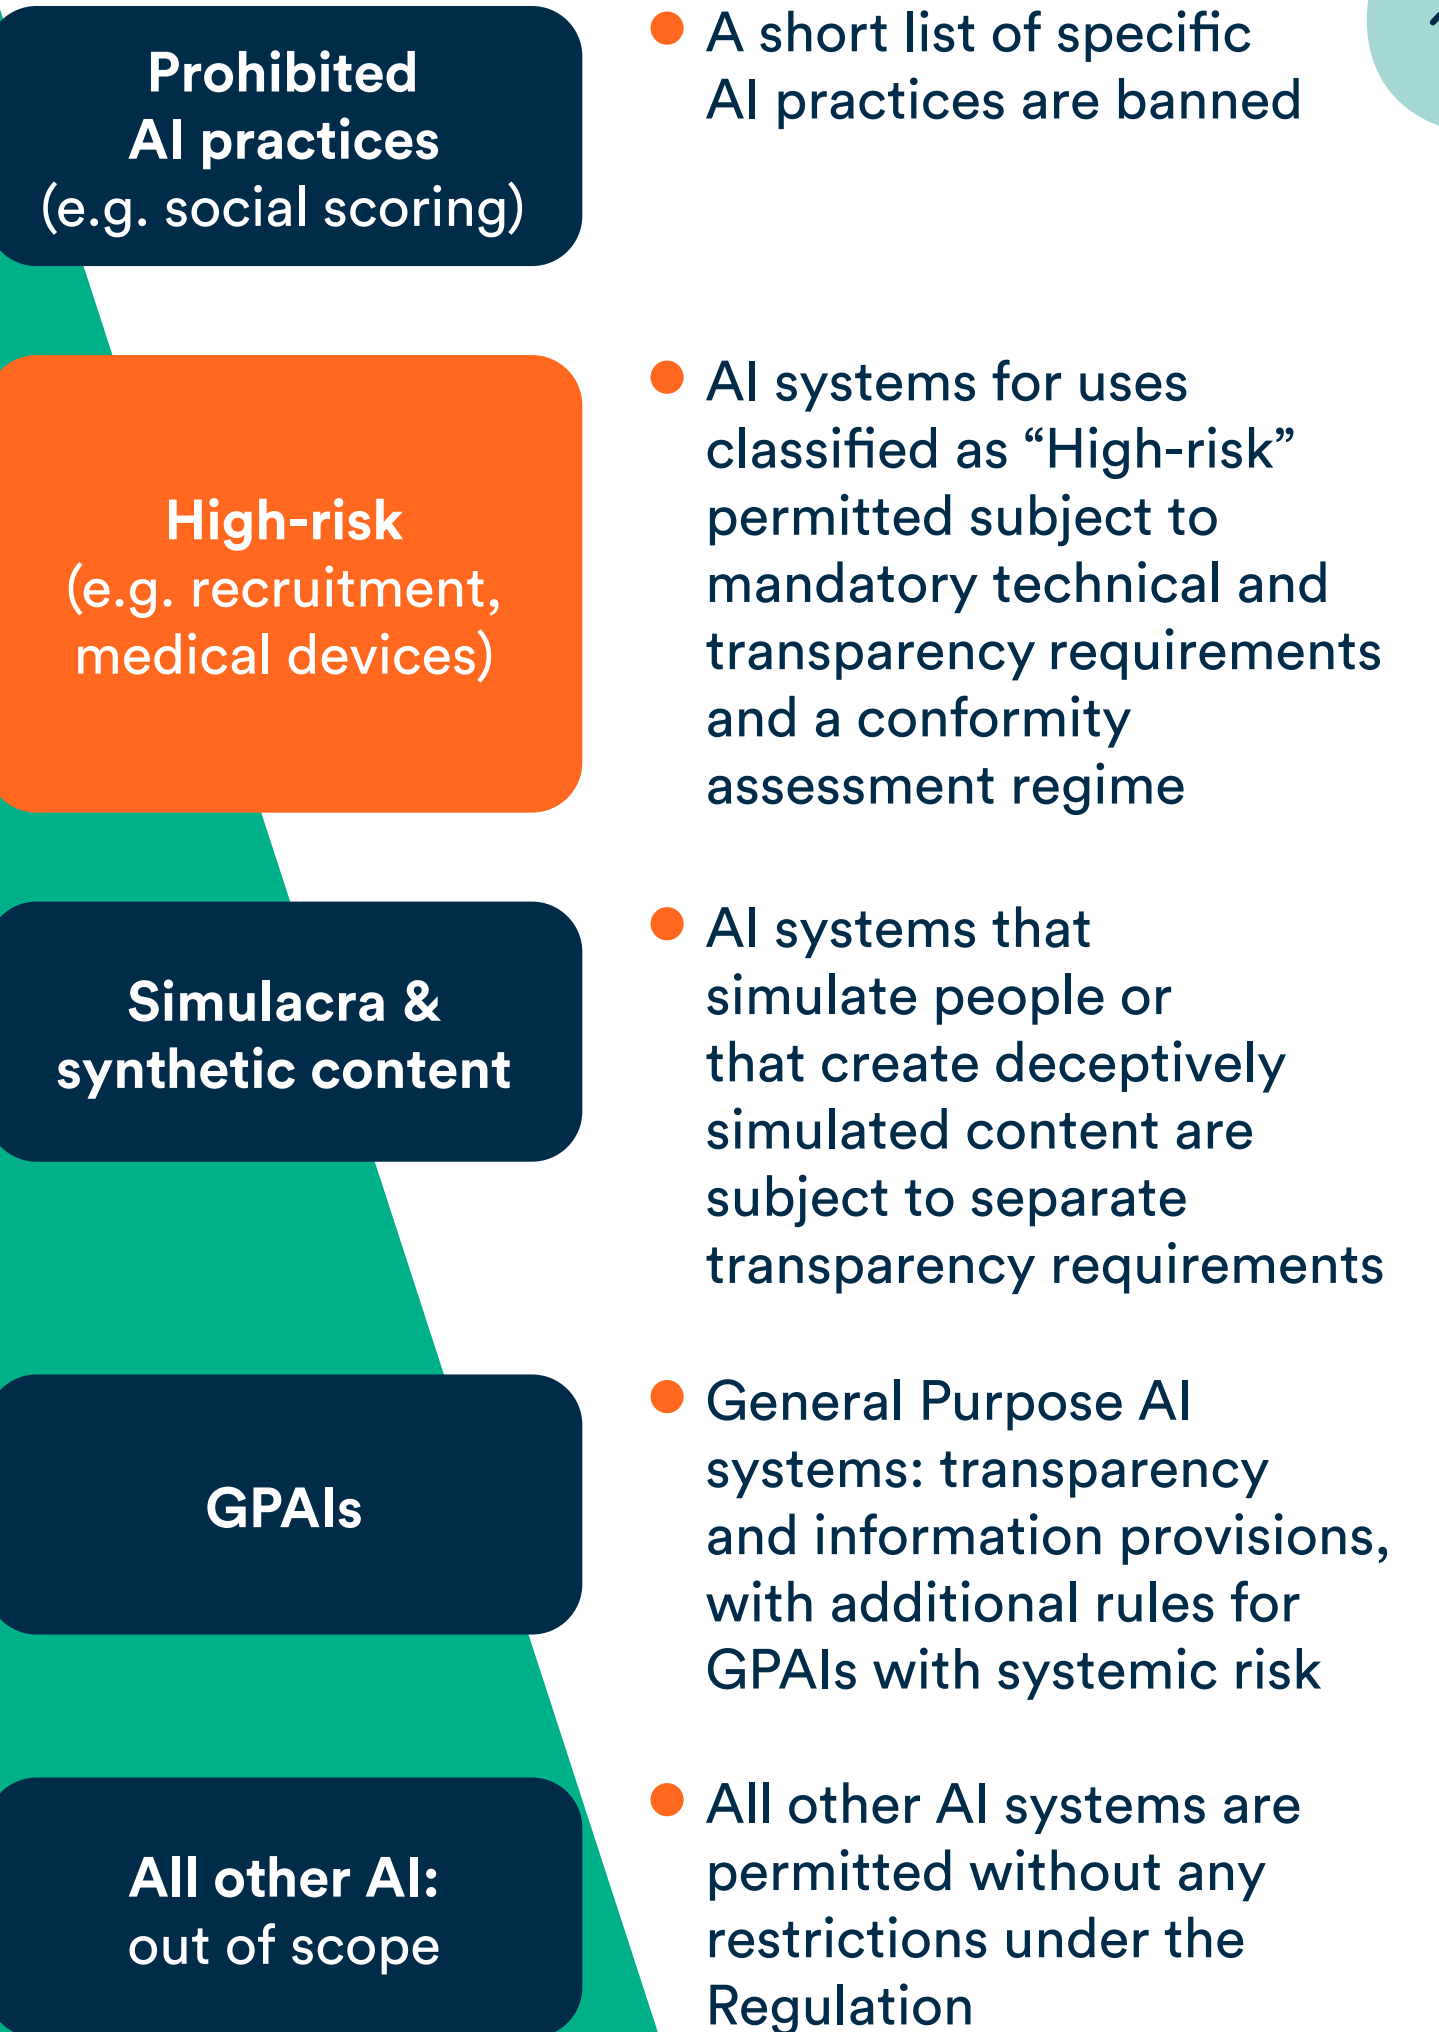
So if your AI system is within scope of the Act, how do you figure out whether it is high-risk or not?

The starting point is to understand that the Act conceptualises “high-risk” in two ways: harm to the health and safety of individuals, and harm to the fundamental rights and freedoms of individuals as enshrined in the EU’s [Charter](#) of Fundamental Rights.

These are two very different types of risk. The potential risks of physical harm to people from AI-powered products such as medical devices, toys and machinery going awry is obvious. The risks to rights and freedoms perhaps less so, but here the focus is on AI systems that influence decisions that may impinge on these rights, particularly in the public sector. For example, the right to asylum, to not be discriminated against and the right to education.

Keeping the distinction in mind between physical harm and harms to fundamental human rights will help you navigate the Act’s rules on high-risk AI.

AI Risk Classifications



How does the high-risk categorisation work?

The Act creates two broad categories of high-risk AI: high risk **products**, and high-risk **use cases**. As a rule of thumb, the risks relating to products are predominantly health and safety risks, and the risks relating to use cases are predominantly risks relating to fundamental rights. The Act formulates each of these in a different way.

High-risk products

The Act deems AI systems as *high-risk* products by referring to a list of EU product safety laws set out in Annex I. The list is split into two sections:



in the first section are product safety laws of a range of products, notably medical devices, machinery, toys and radio equipment;




the second section lists product safety laws relating to forms of transport, mostly planes, trains and automobiles.

The difference between these two sections is crucial. Almost none of the Act applies to the second list covering planes, trains and automobiles, other than some minor provisions including to ensure consistency with the technical requirements of the Act.

The laws in the first section of Annex I are what you should focus on. AI systems in products covered by these laws are classified as high-risk if they meet both of the following criteria:



that the AI system is intended to be a safety component of a product, or is itself a product; and



that the product is required under the relevant law to undergo a third-party conformity assessment prior to it being placed on the market or put into service.

The fact that both criteria must apply is important. Under the laws in question, only a subset of products have to undergo a conformity assessment by a third party. In many cases, the manufacturer is permitted to perform the conformity assessment themselves. If you are developing or supplying these products already, you will know which products have to undergo a conformity assessment by a third party, and which can be self-assessed. We explore the more difficult question of how well the Act integrates with these Annex I laws such as the Medical Devices Regulation and the In-Vitro Diagnostics Medical Device Regulation in [this article](#) – and [see below](#) for the Act's close parallels with, and lessons that can be learned from, the world of medical devices.

High-risk use cases


The Act takes a different approach to classifying *high-risk use cases*. Rather than cross-referring to a list of laws, it refers to Annex III which describes AI systems in specified use cases within specified sectors as high-risk. For example, “critical infrastructure” is a sector, and AI systems used as safety components in the supply of water, gas, heating or electricity are automatically deemed as high-risk uses cases. To take another example, in the “employment” sector, AI systems used to analyse and filter job applications are classified as high-risk.

High-risk use cases



Biometrics

- Remote biometric identification
- Biometric categorisation per protected characteristics
- Emotion recognition



Critical Infrastructure

- Safety component of system used in critical digital infrastructure, road traffic or supply of water, gas, heating and electricity



Education

- Determining access to educational institutions
- Assessments and/or admission tests
- Determining level of education provided
- Cheat detection



Law Enforcement

- Predicting likelihood of person being a victim, assessing evidence, polygraphs and similar
- Assessment re-offending risk
- Profiling for crime-related analytics



Employment

- Recruitment or selection
- Promotion, task allocation and termination
- Evaluating performance and behaviour



Administration of Justice & Democracy

- Assisting a judicial authority in research and application of law
- Influencing outcome of election or voting behaviour



Essential services

- Evaluate eligibility for state benefits and services
- Credit scoring
- Risk assessment for life and health insurance



Migration & Border control

- Verification of travel documents; examination of applications for asylum, visa and resident permits
- Polygraphs and similar for risk-assessment, including a security or health risk

Most of these are in the public sector, but not all are.

The basis on which use cases have been included or excluded is not necessarily intuitive. Emotion recognition systems are included in the biometric sector. Credit scoring is deemed an essential private service but is the only financial services use case that has been included. The list of employment and HR-related use cases is surprisingly long. The bottom line: if you think your sector and/or use case might be in scope, you'll need to read the relevant wording of Annex III carefully to try and discern whether your AI system will be caught.

There is also a set of so-called "filters", that were added to the Act at a late stage, designed to ensure that innocuous deployments of AI systems in the use cases in Annex III are not categorised as high-risk. This means that AI systems intended for the following tasks will not be considered high-risk under Annex III: (a) narrow procedural tasks; (b) improving the result of a completed human activity; (c) analysing human decision-making patterns; and (d) performing preparatory tasks to an assessment relevant for the purpose of a use case. You may notice that the wording of these filters is itself also not immediately clear. What is a "narrow" procedural task, as opposed to a broad one? What does (d) actually mean? It is easy to see the filters themselves becoming contested.

The good news is that the EU Commission has an obligation under the Act to publish guidelines to assist in interpreting high-risk for the purposes of Annex III. The Act states that these guidelines must include a comprehensive list of practical examples of uses cases that are high-risk and not high-risk. The less good news is that the deadline for publishing the guidelines will be in March 2026 (assuming the Act comes into force in August 2024). This will be just six months before most of the Act will apply as law. Whilst this may be too late for developers of AI systems that do not map neatly onto Annex III, it is possible that the Commission's new AI Office will publish the guidelines earlier, or at least provide informal guidance in webinars and other forums in the meantime, perhaps in the context of the [AI Pact](#) initiative.

Conclusion

The rules around high-risk AI products and use cases are complex and rarely intuitive, but understanding the principles behind the Act's approach to high-risk AI should help in applying them to your AI system.



What is my role under the AI Act?

If your AI system might be “high-risk” under the AI Act, the next question to ask is what role(s) and associated responsibilities does my organisation have under the Act? This is critical, because the Act’s obligations relating to high-risk AI systems are determined by the position that the Act assigns to the organisation in the value chain of high-risk AI systems that is created by the Act.

The concept of a value or supply chain of products, conceived as a series of economic operators occupying different tiers in the sequence by which a product proceeds from manufacturer to end customer, comes from the EU’s “New Legislative Framework” (NLF) product safety legislation. But the AI Act takes the idea further by placing obligations not just on entities that “provide” (and, to a far lesser extent, “distribute” or “import”) high-risk AI systems, but also on those that “deploy” them. These two roles, the Provider and the Deployer, are the two most important in the Act.

Once you determine your role, then the Act sets out your obligations in relation to high-risk AI systems in an apparently tidy way. The obligations on Providers are listed in Article 16, and those for Importers, Distributors and Deployers are set out in Articles 23, 24 and 26 respectively. Much of the rest of the Act is focussed on the obligations of Providers of high-risk AI systems, and establishing a framework for enforcement of the Act.

In fact, the potential roles and responsibilities that can accrue to an organisation under the Act go beyond this orderly set of roles in the AI value chain. And there is another important actor hidden away in Article 25, that of a third party supplier of components of a high-risk AI System, whose responsibilities are less clearly defined, but seem likely to have a wide impact in practice.

Roles in the AI value chain

We summarise the main roles as follows:



Provider

An entity that develops (or commissions) an AI system or a general purpose AI model and places it on the market or puts it into service under its own name or trademark, whether for payment or free of charge.



Deployer

An entity using an AI system under its authority, except where the AI system is used in the course of a personal non-professional activity.



Distributor

An entity in the supply chain, other than the provider or the importer, that makes an AI system available on the EU market.



Importer

An entity located or established in the EU that places on the market an AI system that bears the name or trademark of an entity established outside the EU.



Authorised representative

An entity located or established in the EU that has received and accepted a written mandate from a provider of an AI system or a general-purpose AI model to, respectively, perform and carry out on its behalf the obligations and procedures established by the Act.

Many organisations making use of AI will be Deployers

Organisations that use high-risk AI systems will be Deployers, so will be the most common type of role created by the Act and we expect will capture many businesses making use of AI. The obligations on Deployers in Article 26 are not trivial. What kind of technical and organisational measures will Deployers have to take to ensure that AI systems are used in accordance with the Provider's instructions of use? What roles will be given over to the human oversight function required by the Act? How will a Deployer ensure that input data is relevant and sufficiently representative? The optimal outcome would be that a combination of guidance from the EU AI Office and harmonised technical standards will provide practical insights to help answer these questions.

Complex AI supply chain means other roles are harder to determine

The definitions of Provider, Distributor and Importer are all built on the fundamental NLF product safety law concept of the placing on the market of a product, along with the related terms of making available on the market and putting into service. These terms are all defined in the Act, and in many circumstances will be relatively straightforward to apply.

However, given Bristows' long experience in advising on these terms in the context of the EU Medical Device Regulation, especially as that legislation applies to Software as a Medical Device, we anticipate that, in many scenarios, applying these terms to AI systems will present considerable practical challenges. This is partly because the terms were originally defined with physical products in mind and the read-across to software is not necessarily intuitive, and partly because the access and distribution channels in the AI product ecosystem (e.g. API, SaaS, closed v open source) are already far more complex than when the Act was conceived and drafted, both technically and in the sense of the underlying legal/contractual relationships.

Guidance on and examples of roles expected to be needed

Surprisingly, there does not appear to be any specific guidance scheduled for publication on this topic by the AI Office. Article 96(1)(e) obliges the Commission to publish detailed guidance on the relationship between the Act and other EU legislation, including product safety legislation. It would be optimal for, at the least, that guidance to include a section on how to interpret placing on the market in the specific context of the AI Act, accompanied by a long list of detailed examples. Similar guidance exists in the form of the EU Commission's Blue Guide, which is the authoritative reference manual used by practitioners to interpret placing on the market and ancillary terms in the context of the EU's NLF product safety legislation.

However, the Blue Guide does not consider the specific difficulties associated with interpreting these terms in the context of software and AI, because it is primarily intended for use in relation to tangible products which are manufactured, warehoused, shipped and physically supplied/installed. The Blue Guide also does not address the Deployer role, because that role does not exist in other NLF product safety legislation. AI-specific guidance, which takes into account the intangibility of software and all of the different distribution models which are available for AI systems, would be of great value.

Role of Provider is critical to understand

What the Act does usefully do is explain the circumstances in which a Deployer, Importer, Distributor or other third party can become a Provider of a high-risk AI system that is already circulating in the market.

This is explained in Article 25, which we expect will become one of the most heavily cited (and perhaps contested) provisions of the Act. Given the compliance overhead that will fall on Providers of high-risk AI systems, no organisation will want to become one inadvertently. The fact that putting your name or trade mark on a high-risk AI system that is already on the market is enough for your organisation to be deemed the Provider of that AI system is a key takeaway here.

Article 25 also introduces another role in the value chain: the supplier of an AI system, tools, services, components or processes that are used or integrated in a high-risk AI system. Because it is not neatly defined, and only appears in Article 25(4), this component supplier role has not attracted much attention or commentary, but our view is that it could catch a broad swathe of vendors who otherwise consider themselves out-of-scope of the Act.

The obligation is for a written contract to be used between Providers of high-risk AI systems and these suppliers to facilitate the Provider's compliance with the Act. The AI Office is to produce a set of voluntary model contractual terms for this purpose.

It is interesting to consider this in light of the so-called LLM orchestration stack that has emerged over the last 18 months in relation to foundation models, with a variety of highly specialised service providers (e.g. vector databases, open-source LLM platforms) now available to assist in developing a LLM-based product. Given how widely scoped the component supplier role is under Article 25(4), it is easy to see the AI Office's voluntary model terms being very widely adopted if (as seems inevitable) a similar technology/vendor stack develops around high-risk AI systems, or at least some of the more popular use cases.

Overlaps and closing thought

Article 25 raises the possibility of AI systems developed by a third party, being integrated into a Provider's high-risk AI system. The same is true of general purpose AI models under Article 54, which could be a component of an AI system (whether high-risk or not), or become a standalone high-risk AI system.

Similarly, the transparency provisions in Article 50 (relating to lifelike interactive AI, deepfake content, etc.) will apply to any AI system that has the relevant functionality, which could be a high-risk AI system and/or a general purpose AI model.

So, whilst the value chain of Provider, Importer, Distributor and Deployer of high-risk AI systems is still the core set of roles around which the Act is organised, they are not the whole story, and we are only beginning the journey of understanding how to interpret their obligations under the AI Act.



Are data protection governance frameworks a good model for AI governance?

What feels really British but isn't? "Almost everything in the *British Museum*" goes the old gag. Could a similar analogy be made about the role of data protection governance frameworks as a model for shaping AI governance? Does data protection feel as if it might be a perfect fit for shaping AI governance, and actually isn't?

Reasons to be cheerful ("DP governance = AI governance")

There is a school of thought that for effective AI governance, organisations could do worse than start with the data protection governance frameworks that many, particularly in Europe, already have in place. Here are some of the reasons we commonly hear in favour of this:



Many organisations in Europe now have a **data protection governance framework** and a fast growing number operating further afield do too. In Europe a few of these pre-existed the GDPR, but many were created for GDPR and so many organisations have experience in living memory of setting up such frameworks. Why waste effort that has already been invested in setting up and maintaining these structures when instead they could be adapted for AI governance?



The features of typical data protection compliance and governance programmes (at least in the EU and UK) generally cover a lot of the ground needed for an effective AI governance programme. For example, a **data inventory** (including mapping data flows) looks broadly like the sort of inventory for which the equivalent is needed in relation to AI systems.



Similarly, the **third party supplier contract due diligence** exercises many organisations conducted during GDPR implementation sound useful for discovering whether your existing suppliers are already using AI to deliver services to you but may have neglected to share this with you. Other features of data protection governance only serve to reinforce this view.



Both the fields of data protection and AI are "team sports" requiring the **participation of stakeholders** with different qualifications, and also different organisational responsibilities, to be effective. Privacy professionals are already adept at leading multi-disciplinary teams.



Many of the **skill sets necessary for effective governance** look essentially similar in both fields. Take, for example, the concept of ongoing monitoring under the AI Act; doesn't this look in effect loosely similar in intended outcomes to the concept in the GDPR of "accountability"?

Reasons to think twice

Where are the limits to the arguments above? How well do they withstand scrutiny?

The primary purpose of the Act is to prevent risks caused by “high-risk” AI systems (read “[Is my AI high-risk under the AI Act?](#)” article). For that reason, we confine ourselves below to assessing only how high-risk AI systems measure up against these arguments.

At the heart of the AI Act is the Title III regime, which governs AI systems that are deemed to pose a “high-risk” as Recital 43 and Article 7(2) make clear, to “health, safety and fundamental rights”. As noted in this series ([see article](#)), the twin pillars of the Act’s approach to its concept of “high-risk” AI systems are, on the one hand, the health and safety of individuals, and on the other, the fundamental rights and freedoms of individuals as enshrined in the EU’s Charter of Fundamental Rights.

This regime is based on a common EU approach, which has been around for decades, to regulate products where safety is of particular importance – medical devices and lifts to name two examples.

Under this approach, the manufacturer must establish the safety of regulated products through conformity assessments of the products against certain essential statutory requirements. This must be done before the products can be placed on the market (first commercialised). Thereafter, products may be marked “CE”, allowing their marketing and distribution across the EU. However, it should be stressed that this is not the end of the process and such products are subject to continuous monitoring and vigilance obligations.

Providers of high-risk AI systems, on whom most of the AI Act’s essential requirements fall (see Chapter 2 of the Act), must create and operate a quality management system for the AI systems they have developed. The Act’s requirements for the quality management system are set out in Article 9.

If this regime sounds hardly like the legislative regime for data protection in the EU at all, that is because it is not. There is little similarity between these core features of the Act in relation to high-risk AI systems and the EU legislative approach to data protection as embodied in the GDPR.


Delving further, conformity assessment as a regulatory model depends on standardisation organisations and notified bodies for its efficacy. Broadly, providers that follow a standard developed by one of the European Standardisation Organisations do not need to interpret the essential requirements of the legislation but will instead simply be able to follow the relevant standard. The capacity problems in the EU’s system of regulation for this approach are ably described in the [following article](#), but that is another matter. The GDPR’s attempt to kickstart a market in mechanisms loosely akin to such standards, that is, certification schemes, certification bodies, and codes of conduct for data protection has been one of its notable failures. Six years on from GDPR implementation, the rate of adoption by data controllers of such certification schemes and codes of conduct remains underwhelming and undoubtedly a disappointment to policymakers and market participants alike.

Conclusion


So why is the school of thought advocating data protection governance as a model for AI systems governance getting the traction it is? Notwithstanding its shaky assumptions about the legislative scheme of the AI Act (i.e. very different to that of the GDPR), it seems that, at a high level at least, some aspects of governance programmes probably do reflect business processes that do not change that much whatever their subject matter. It seems that, for now at least, organisations are quick to see the opportunity to increase the return on investment in frameworks that they have already developed (for GDPR) by recycling them for AI systems:



It is realistic to conclude that **AI governance cannot safely be parked with one function or role** (e.g. the CIO) and left to thrive safely there without input and oversight from, in all likelihood, several other functions. The roles of the DPO's office in data protection and the CIO, CISO, Legal, Compliance & Ethics, Internal Audit and the business do not seem so over-engineered after all.



It does seem sensible for an organisation to **make an inventory early on in its governance process** of AI systems that: (a) it is already using internally; (b) it wishes to deploy in the near or medium term; and (c) that its suppliers are already using to deliver services to it (sometimes without having informed the client). Such inventories are, superficially at least, not unlike the data mapping and contract inventories maintained under GDPR.



Documentation: perhaps this is one person's icing on the cake and another's "killer app". What percentage of good governance is attributable simply to **documenting your processes, your controls, your "guardrails" and your mission statement**? Ask any data protection specialist and you will find documentation is crucial. The requirements of the AI Act, as we wait with bated breath for a deluge of guidance from the Commission, ENISA and other bodies over the coming 12 to 18 months, are very similar in this regard.



The EU AI Act as product safety legislation – lessons from the world of medical devices

If an AI system is classified as “high-risk”, the AI Act imposes significant requirements on the provider in order for it to place the system on the market and then on an ongoing basis. These requirements mean that, in many ways, the AI Act amounts to a piece of product safety legislation. As such, organisations in those sectors well-versed in placing safety-critical products on the market are better prepared to comply with the product safety aspects of the Act than others. One such sector is life sciences, where there are very close parallels between the product safety features of medical device legislation and those of the AI Act. Our specialist life sciences regulatory team regularly advises clients on these matters, from risk classification, product liability, the need for conformity assessments, CE-marking and post-market surveillance, all of which are mirrored in the AI Act.

Below, we explore a few of those parallels and offer some potential lessons from medical device legislation for those seeking to comply with the product safety features of the AI Act, in the life sciences sector and beyond.

Medical device law – quick background

In 2017, the EU adopted twin new regulations: the Medical Devices Regulation and the In Vitro Diagnostics Medical Devices Regulation. These represented significant changes to the relatively well-established regulatory frameworks for medical devices and in-vitro diagnostic medical devices.


These new regulatory frameworks represented the biggest change in thirty years and imposed a number of more stringent requirements on “Manufacturers” (the entities with primary regulatory responsibility for a device). However, the new requirements were not reserved to “Manufacturers”: new requirements were also imposed on other economic operators involved in the design, development and supply of components of medical devices.

This caused several issues for organisations involved in the supply chain, many of which seem set to be repeated by the EU AI Act. In fact, the AI Act is even more ambitious than the new economic operator requirements in that it also seeks to regulate entities that deploy AI systems and in certain instances, suppliers of components to high-risk AI systems. This will result in significant additional compliance obligations in the supply chain.



No “grandfathering”

Every existing approved (CE Marked) device needed to undergo a new Conformity Assessment under the new, more stringent requirements. This created an instance backlog as there was a rush to re-certify existing devices.



Even more revolutionary for IVDs

Under the previous Directive only around 10% of IVDs were required to undergo a Conformity Assessment involving a Notified Body. Under the IVDR this reversed and around 90% of IVDs were required to undergo a Conformity Assessment involving a Notified Body.



New clarifications


As is always the case with a new regulatory frameworks such as these, the European Commission needed to publish a numerous guidance documents and standards to clarify the new requirements. Unfortunately, these were (and continue to be) seriously delayed.



More data required and scrutiny

Manufacturers were required to gather and analyse more clinical data even for existing well-established devices. As such it became necessary to conduct additional clinical studies to gather clinical data to establish the safety and performance of devices.

Finally, certain higher risk innovative devices now need to undergo an additional scrutiny procedure after having successfully concluded the Conformity Assessment. This was memorably described as Putting the No in InNOvation.



Higher standards and up-classification

New more stringent requirements were imposed and a number of medical devices were “up-classified”. Most notably, the vast majority of software medical devices moved from being self-certified to requiring a Conformity Assessment undertaken by a Notified Body. Precisely because such a small fraction of software medical devices previously required the involvement of a Notified Body, they had little experience with software medical devices. This presented a particular and continuing challenge.



Existing Notified Bodies needed to be re-certified and upskill

Every existing Notified Body had to be re-certified under the new enhanced standards. This distracted Notified Bodies significantly at precisely the point in time when they were in highest demand by existing customers, let alone developers if manufacturers of new devices. These all generated a huge backlog of devices waiting to be re-certified. In turn, this delayed the conformity assessments required for new devices.



Additional infrastructure: EUDAMED - one database to rule them all

Finally, a central pillar of the new frameworks is a new database called EUDAMED intended to collate and process information regarding devices, economic operators, notified bodies, clinical studies, vigilance and surveillance.

Seven years after the new regulations were adopted, EUDAMED has still not been completed. In March 2024, the European Commission further amended the regulatory MDR and IVDR to give more time. It now appears likely that EUDAMED will not be fully functional until 2026 or 2027.

Tight deadlines got extended and now proposed changes

While all of these changes were well-intentioned and improved the regulatory framework, the implementation caused enormous difficulties and the withdrawal of products. These difficulties resulted in delays to the implementation of the new frameworks, extended transitional periods and now a proposal for a slew of amendments.

In our view, the main causes of the disruption associated with MDR and IVDR were:

- Unrealistic implementation timetables.
- A shortage of Notified Body capacity – especially as regards software and in vitro diagnostics.
- A lack of guidance and delays issuing the guidance. In all honesty, when some of the guidance was finally published, it was so ambiguous that the rushed guidance caused difficulties.
- Inconsistent interpretations and approaches adopted by different member states and different notified bodies.
- Overly optimistic timetables to build and deploy a database.

Likely to repeat each of these

Unfortunately, each and every one of these missteps appear likely to be replicated with the adoption of the AI Act particularly for “high-risk” AI systems, which represent the primary portion of the AI Act. We can already see:

- A disastrous shortage of competence and capacity at Notified Bodies.
- Virtually none of the guidance necessary to implement the AI Act has been published
- We still do not have the new central regulator or, in many cases, the national competent authority.

This is particularly disconcerting given that the requirements in the AI Act require all stakeholders to come to grips with totally new concepts like bias management and governance.

Further, frustratingly, there are a number of significant challenges reconciling the AI Act with the requirements under the MDR and the IVDR. This is crucial as the two regulatory frameworks are intended to operate in an interconnected manner. By way of example, the language in the AI Act is inconsistent with the accepted terminology for Conformity Assessment of existing products like medical devices. Worse, there are now some instances of serious unintentional conflicts between the requirements of the AI Act and the MDR and IVDR, whereby conducting an authorised clinical study in accordance with the MDR or IVDR might constitute an offence under the AI Act.



Copyright issues under the EU AI Act

Intellectual property is one of the main issues for everyone involved in the burgeoning AI ecosystem. The large generative AI model developers, have been coming under scrutiny, and in some cases legal claims, regarding their use of copyright content to train their models. Simultaneously, organisations using large models are concerned about whether intellectual property subsists in content generated by these tools.

So what does the AI Act have to say about intellectual property? The Act was not, at least at the beginning of its journey, intended to regulate copyright. However, as the proposal advanced, large language models developed exponentially, and pressure grew to address some of the copyright concerns they gave rise to.

So where did we end up? There are three aspects of the approved AI Act which are particularly relevant to copyright.



Text and data mining and training general purpose AI models

Under Article 4(3) of the Digital Single Market Directive, it is permissible to make copies of lawfully accessible works for the purposes of text and data mining, including commercial purposes. It is, however, open to copyright holders to “opt out” of this exception. To do so, the copyright holder is required to expressly reserve the right of text and data mining to themselves. Exactly how the copyright holder does this is less clear. The DSM Directive suggests that it should be in an “appropriate manner”, such as using machine readable means.

Prior to the final version of the AI Act, there had been some uncertainty and debate about whether the text and data mining exception could apply to acts of copying copyright works in order to train general purpose AI models. This is for two reasons.

Firstly, text and data mining on the one hand, and AI model training on the other, are not the same activity. Text and data mining (defined in the DSM Directive as “any automated analytical technique aimed at analysing text and data in digital form in order to generate information which includes but is not limited to patterns, trends and correlations”) typically involves scraping data, extracting relevant data and pre-processing it. AI model training generally involves taking a selected model type, applying it to a training data set and then making the necessary adjustments to fine tune the model. Perhaps a sensible way of looking at this issue is to think of text and data mining as a necessary prelude to AI model training.

Secondly, it has been questioned whether or not the EU legislators had the development of general purpose AI models in mind when framing the text and data mining exception in Article 4(3) of the DSM Directive (which was finalised in 2019).

The recitals of the AI Act appear to dispel the uncertainty.

Recital 105 clarifies that text and data mining techniques may be used in order to retrieve and analyse the vast amounts of text, images, videos, data etc. that are required to train general purpose AI models, and that this typically requires the authorisation of the copyright holder, unless a copyright exception applies. So the recital expressly acknowledges that there is a nexus between the use of

copyright works and training general purpose AI models. Recital 106 also expressly references Article 4(3) of the DSM Directive by requiring providers of general purpose AI models to put in place a policy to ensure that they comply with any reservation of rights under Article 4(3), and this is reflected in Article 53.1(c) of the AI Act.

Accordingly, whether or not the EU legislators had the development of general purpose AI models in mind when framing the text and data mining exception in Article 4(3), the AI Act now appears to put this issue to bed.



Transparency in relation to the use of copyright works in training materials

One of the biggest challenges for copyright holders with general purpose AI models that the AI Act seeks to address is transparency. The challenge is this: if copyright holders do not know whether their copyright works have been used to train a general purpose AI model, enforcing their copyright against the developer of that model is much more challenging.

Article 53.1(d) of the AI Act imposes an obligation on the developer of a general purpose AI model to draw up and make publicly available “a sufficiently detailed summary about the content used for training of the general purpose AI model” in accordance with a template provided by

the EU AI Office. The nature of this “sufficiently detailed summary” is elaborated upon in recital 107. It should “be generally comprehensive in its scope instead of technically detailed to facilitate parties with legitimate interests, including copyright holders, to exercise and enforce their rights...for example by listing the main data collections or sets that went into training the model, such as large private or public databases or data archives, and by providing a narrative explanation about other data sources used.” There had been some concerns with earlier proposals that too much detail would be required from general purpose AI developers, making compliance impossible and the provision unworkable.

As it now stands, the obligation should – subject to the EU AI Office template once issued - enable developers of general purpose AI models to provide a relatively high level explanation of their data sources that enable copyright holders to determine whether they are “lawfully accessible” data sources, that include their copyright works. It seems to assume, however, that the copyright holder will already know whether their work is included in a particular data source, which may not be the case.

There is a balancing act going on here. Copyright holders want as much detail as possible to make it easier for them to enforce their rights. Conversely, AI developers want reassurance that they will not face a slew of lawsuits that will make their operations unviable.



Long arm jurisdiction for copyright

Article 53.1(c) and recital 106 of the AI Act make it clear that providers who place general purpose AI models on the EU market are required to ensure compliance with the AI Act and implement a policy for doing so, including text and data mining exception in Article 4(3) of the DSM Directive (by using state of the art technologies).

Recital 106 goes on to state that “Any provider placing a general-purpose AI model on the Union market should comply with this obligation, regardless of the jurisdiction in which the copyright-relevant acts underpinning the training of those general-purpose AI models take place”. This is very much a “long arm” jurisdiction. It applies the compliance obligation to providers of general purpose AI models who are located outside of the EU and it does not matter where the training of the model took place or what the copyright laws of those countries are. The justification for this approach is to ensure that no provider of a general purpose AI model gains a competitive advantage within the EU by applying a lower standard of copyright.

Not so long ago, the EU legislators introduced a long arm jurisdiction under the GDPR. So it is perhaps unsurprising to see a similar approach taken in the AI Act. Of course, it may be desirable to have a uniform standard of copyright concerning general purpose AI models and the EU is clearly looking to seize the initiative in the AI Act. However, it does stretch the limits of international comity and is one of the more controversial copyright aspects of the AI Act.



What employers need to know about the EU AI Act

Uses of AI systems in an employment context are to be tightly regulated by the AI Act given the potentially significant impact their use could have on a person's career prospects and ability to earn a living. There is a real risk that AI systems used in recruitment may perpetuate historic biases, e.g. against women or people of certain races, and using AI systems to monitor an individual's employment performance risks interfering with their fundamental rights and privacy. Note that, in common with the provisions governing high-risk AI systems, the key employment-related provisions will not apply for another two years from the date the Act enters into force.



Why does the AI Act matter to UK employers now that we have left the EU?

Whilst the AI Act may not be relevant to all UK employers, those who have a European footprint and intend to utilise the same AI systems across all territories may well decide to put in place similar arrangements in the UK to those required in EU states. It is expected that the EU's approach to AI will set a minimum global standard for regulation.

Further, the AI Act applies if the output of an AI system is to be used within the EU, so any UK-based employers who target, or accept applications from, EU-based candidates will need to comply with the AI Act to the extent they use AI as part of their recruitment process.

What is a high-risk AI system?

As described [above](#), the AI Act identifies certain AI uses as "high-risk", which means that they pose a significant threat to a person's health, safety, or fundamental rights, and are therefore subject to stricter regulation. Various employment-related uses of AI are deemed high-risk. They fall broadly into two categories:

-  Recruitment uses: targeted job adverts, screening of applications and evaluations/selection of candidates.
-  Uses within an ongoing employment relationship: performance evaluation, work allocation on the basis of behavioural or personal traits, and promotion and termination decisions.



Are there any exceptions?

If an AI system does not “pose a significant risk of harm, to the health, safety or fundamental rights of natural persons” then it will not be deemed high-risk. The AI Act gives four specific examples of when a system will fall within this exception/derogation, all of which could potentially be relevant to HR-related AI systems. For example, an AI system intended to improve the result of a previously completed human activity, or an AI system intended to carry out a narrow procedural task.

Further guidelines on high-risk AI systems are to be published by the Commission 18 months from entry into force of the AI Act. These will include practical examples of what would and wouldn't be considered a high-risk system, so this will hopefully provide clarity on what will fall within the high-risk derogation.

What does this mean for employers using or considering the use of AI systems as part of their HR processes?

Most of the obligations relating to high-risk AI systems fall on providers, who are those that develop AI systems or have them developed in their name. In summary, providers of AI systems are required to design systems that:

- have appropriate risk management systems in place;
- that are developed using high-quality data sets;
- for which they can provide detailed technical documentation;
- that automatically keep adequate records;
- are transparent;
- allow effective human oversight; and
- are accurate, robust and secure.

Most employers will be “deployers” for the purpose of the AI Act. The key obligations for deployers of high-risk AI systems are:

- ensuring compliance with the AI system's instructions;
- assigning human oversight to competent individuals who have the necessary training, authority and support;

- monitoring the operation of the AI system and informing the system provider/distributor and, where relevant, the market surveillance authority of identified risks and serious incidents;
- retaining automatically generated logs if under their control;
- informing worker representatives that workers will be subject to the use of a high-risk AI system and informing individuals where an AI system is used to make or assist in making decisions about them; and
- co-operating with the relevant competent authorities in any action relating to the AI system.

Is there anything else employers should be aware of?

Whilst the majority of HR AI use cases will fall within the high-risk category, there are also certain uses that are entirely prohibited. The prohibition most relevant to employers is the restriction on using AI systems to infer emotions in the workplace. Certain AI-powered video interview software already on the market that analyses a job candidate's facial expression may well fall within this category.

Keep an eye out for our future updates on the AI Act as further guidance is published by the Commission.

Related content...

Podcasts: Tune in to the latest insights, on the go

Listeners will hear our industry-leading experts discuss a wide range of topics and obtain a deep understanding into the technology sector trends, perspectives and solutions.



[The Roadmap](#)

Unpacking various technologies or trends such as transformational projects, market practices, or regulatory changes.

[Click to listen](#)



[Legitimately Interesting](#)

Topical commentary on current data protection & privacy issues and developments.

[Click to listen](#)

Articles: Bristows' SnippiTs

Read our expert views on the latest court decisions for the technology sector and beyond.



[Bristows' SnippiTs](#)

[Click to read](#)



Subscribe for frequent updates!

Read our articles, blog posts and other commentary to keep up with the latest legal and regulatory developments relating to AI and machine learning.

[Click to subscribe](#)





Bristows LLP
100 Victoria Embankment
London EC4Y 0DH
T +44 20 7400 8000

Bristows LLP
Avenue des Arts 56
1000 Bruxelles
Belgium
T +32 2 801 1391

Bristows (Ireland) LLP
18 - 20 Merrion St Upper
Dublin 2 D02 XH98
Ireland
T +35312707755

Bristows.com

Bristows LLP is a firm of solicitors and operates as a limited liability partnership governed by English law (registered number OC358808). All references to Bristows are references to Bristows LLP or to Bristows (Ireland) LLP. Bristows LLP is authorised and regulated by the Solicitors Regulation Authority (SRA Number 591711)

Bristows (Ireland) LLP is an Irish partnership of solicitors registered with the Law Society of Ireland (firm number 1262610) and authorised by the Legal Services Regulatory Authority in Ireland to operate as a limited liability partnership (LLP) under the Legal Services Regulation Act 2015 (registered number 1262610) registered with the Law Society of Ireland. The partners of Bristows (Ireland) LLP are all admitted to practise as Irish solicitors.