NORTON ROSE FULBRIGHT

Artificial Intelligence: Foundations for good AI governance

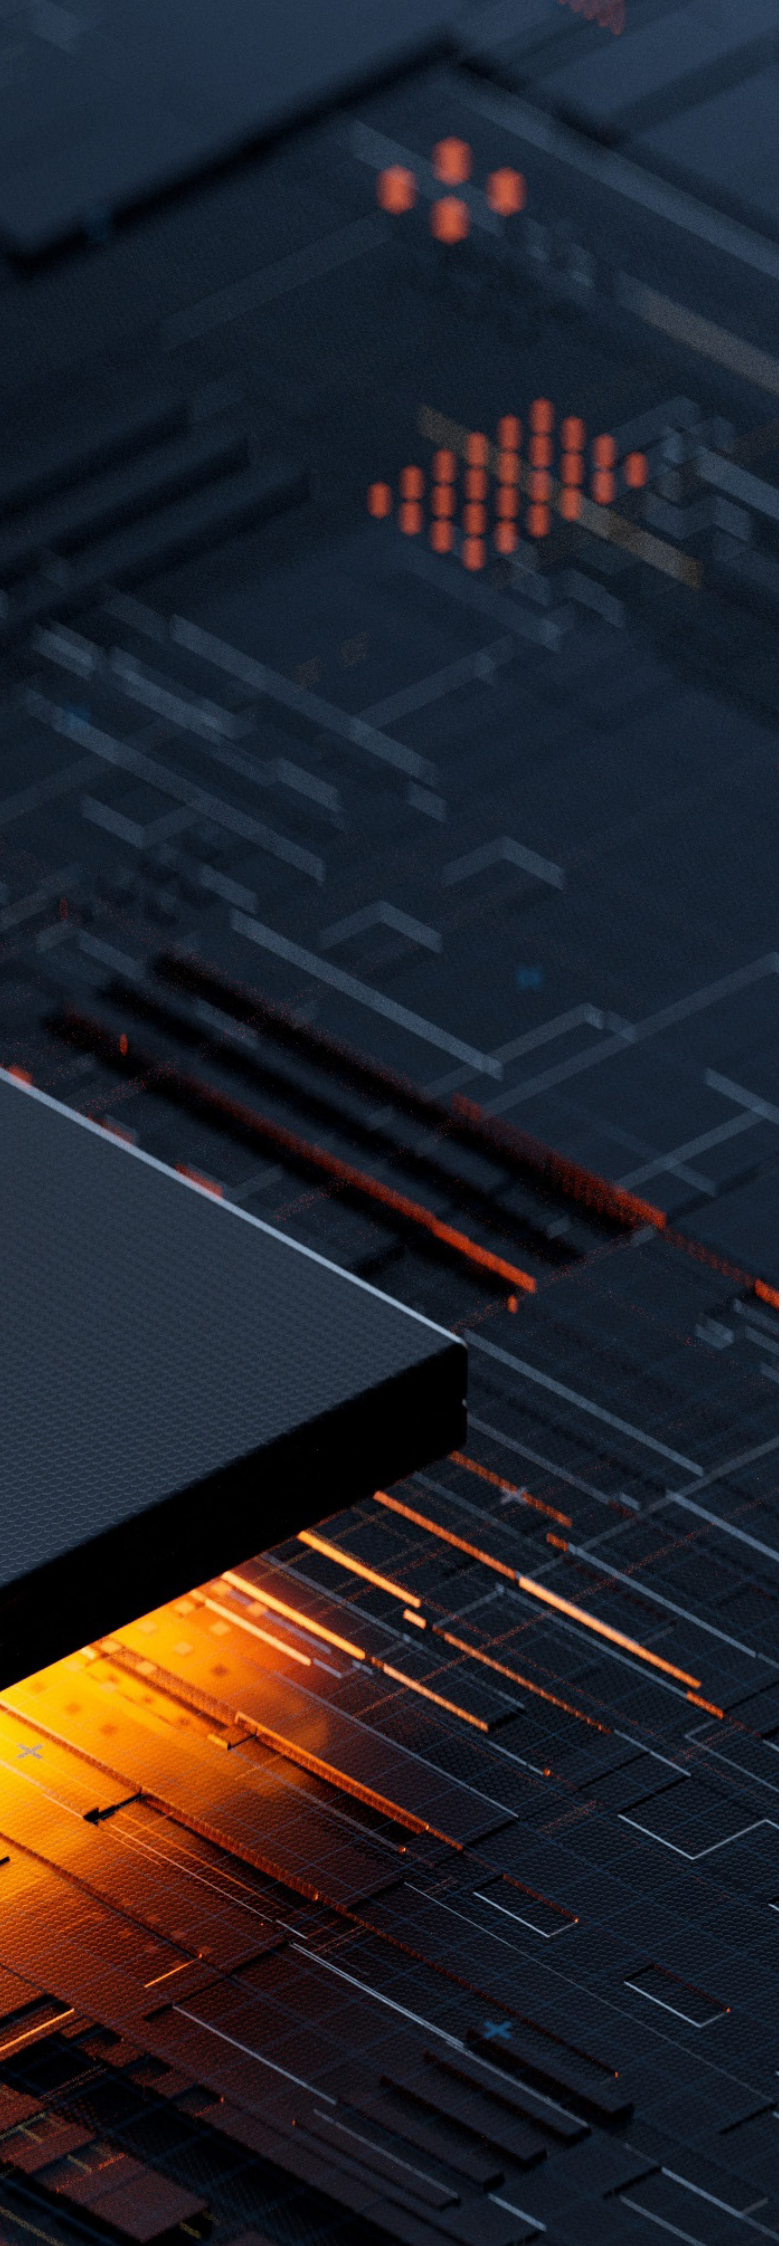

Find out more
Scan the QR code with your smart device for more information.

Contents

Context

Organisations are increasingly using artificial intelligence (AI) to capitalise on the potential of its transformative power. From virtual assistants, to systems that make recommendations for hiring, to automating decisions for credit, to behavioural analysis to better serve a customer segment (to name only a few), AI is being integrated across the value chain in all business domains.
The features of many AI systems provides novel challenges to governance, including issues of bias, discrimination, lack of transparency, and the need to explain and justify decision processes.

Good AI governance is essential to mitigate against material AI risk.

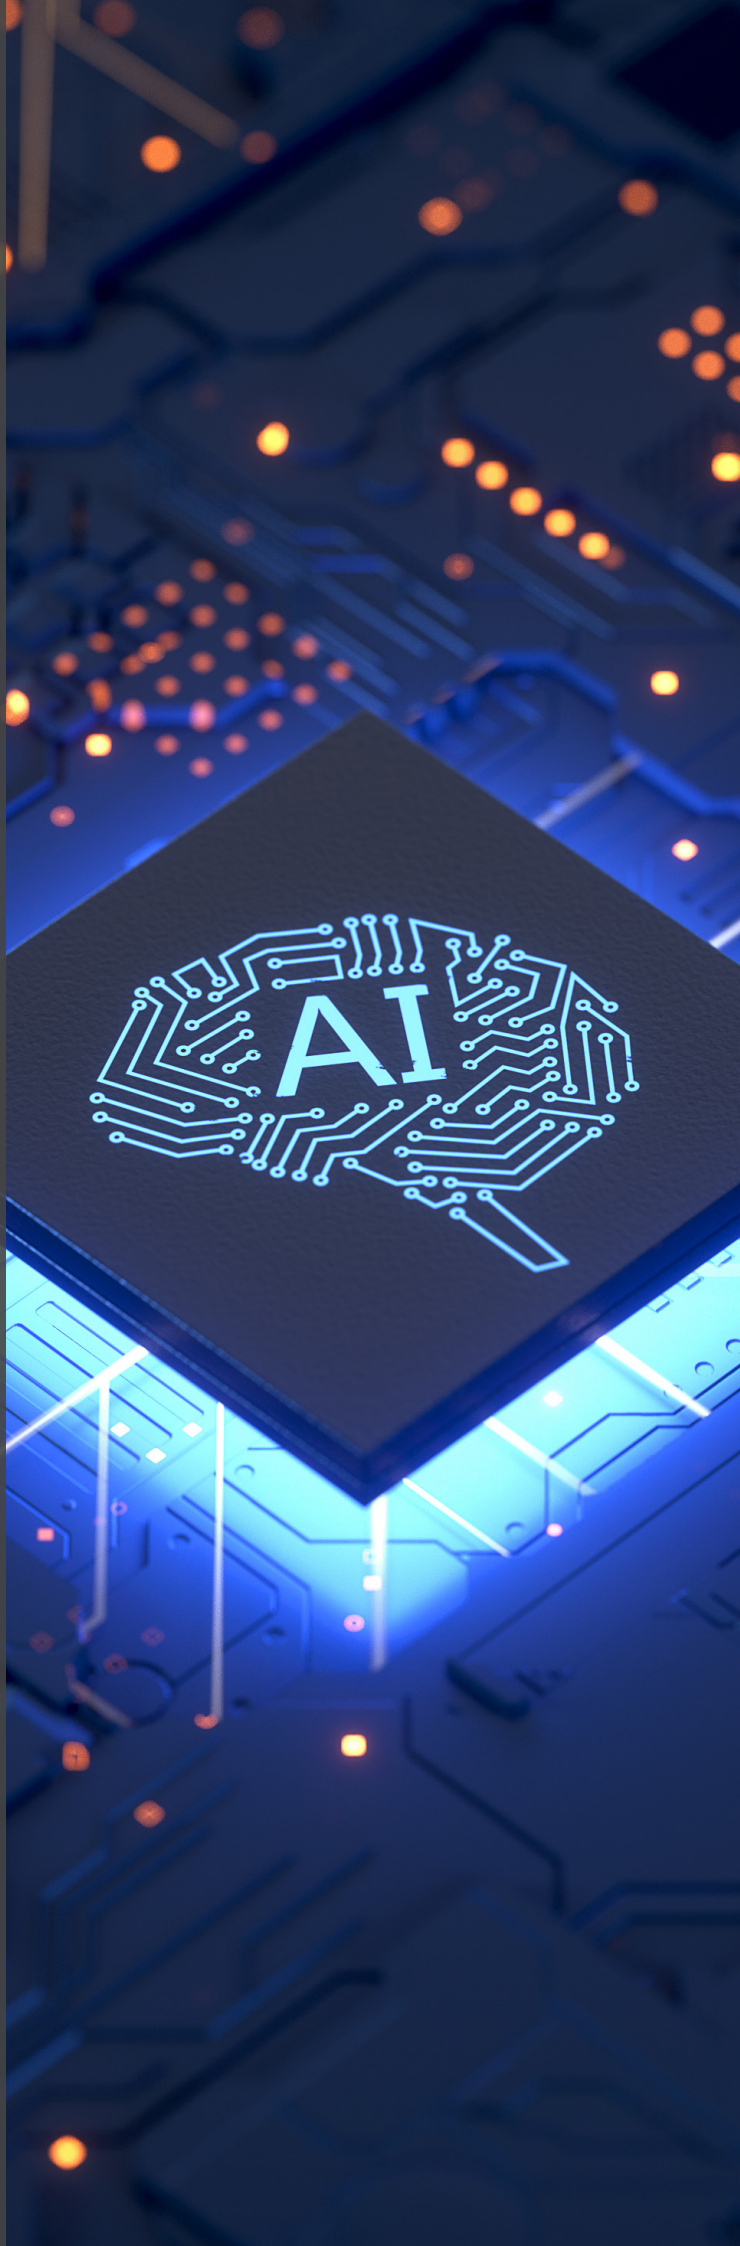

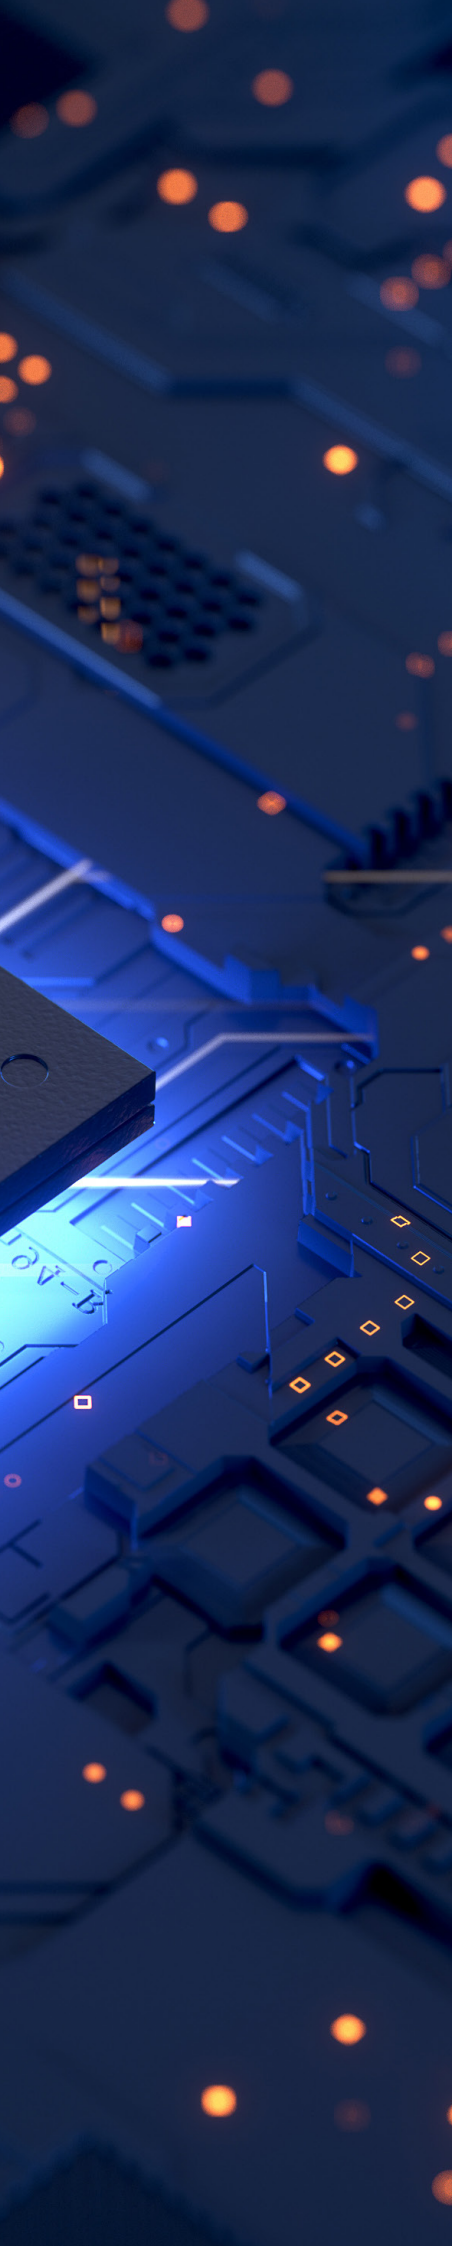

Challenges

Organisations face several key challenges in designing and implementing AI governance frameworks.

- Whilst AI can bring significant upside, AI technologies and systems also bring novel issues and new risks that must be considered and mitigated. This is not always a straightforward task for large organisations. AI can span multiple use cases each of which brings different potential benefits and harms. It is worth highlighting that some simple automation techniques can pose high risk if used incorrectly or inappropriately.
- An organisation is also likely to be exposed to AI risk from multiple dimensions – they might develop AI solutions in-house or in a collaborative venture with a partner, they might procure technology that incorporates AI from a third party, or AI might be used in services an organisation outsources to a third-party provider.
- AI systems and uses must be appropriately risk assessed throughout the lifecycle of the AI system, to protect against harm. It can be difficult to effectively identify, assess, and mitigate the risks of harms of multiple AI technologies, systems and their uses proportionate to the risk they present.
  - AI raises new ethical issues about how we trust AI systems (which may be opaque or autonomous or may change themselves after they are deployed), the way we use data (where it comes from, if it can be used or re-used), and the potential effects of bias and discrimination when AI makes decisions.
  - Potential risks may arise from training or input data which is biased or incomplete, from AI algorithms or techniques which themselves introduce bias or discrimination, from deployed AI systems or overall business processes which are unfair in their treatment of different categories of people, and from AI algorithms or overall processes which are not transparent or not able to be readily explained or justified.
  - As organisations deploy increasing numbers of AI systems across their businesses, we will find that the output data of one system will be the input or the training data of others. In these eco-systems of AI applications, bias or discrimination may percolate through system after system, causing potential harms well beyond the original source.

For reasons such as these, AI has drawn concern from regulators. There is now a clear trend toward increased regulation to ensure its use is safe and ethical.

EU AI Act

The European Union has proposed the AI Act, the first significant law regulating AI development, distribution, training and use by a major regulator (there are a number of less comprehensive proposals in other jurisdictions too). The Act is expected to become law during 2023, with a 24-month transitional period (that is, requiring full compliance some time in 2025). This means organisations don't have long to understand, develop and implement necessary governance structures.

The Act takes a risk-based approach. It prohibits use of some AI systems and practises and places onerous obligations on providers and users of AI systems categorised as being high risk. Significantly, such obligations also apply where a business white-labels a high risk AI system as its own when it has been provided, or operated on its behalf, by another business (a common scenario).

Other regulatory trends

Along with explicit AI regulation, regulators – such as data protection, financial services and consumer protection regulators – are focussing more and more on the need for good AI practices, expanding their existing powers to protect the public against AI risk.

Key challenges

Understanding

Emerging requirements

Best practice for AI governance is still emerging (with no common set of standards), making understanding what frameworks and controls to implement for your organisation a complex task.

Unclear definitions

The definitions of what AI systems, technologies and uses are in scope of different laws and regulations can be unclear. Many best practises are still evolving. This can make implementing detailed processes difficult, particularly those that require enhanced or specialist attention.

Visibility

Silos

Organisations often have disparate AI systems and uses spread across multiple business units and jurisdictions. Organisational silos can result in missing or incomplete information being captured, or inconsistent processes being applied, leading to compliance gaps.

Data

A lack of consistency of data capture and knowledge of which information is needed can lead to a lack of visibility of AI system risk across the organisation, that can make effective oversight a challenge.

Strategy

Insights

Without consolidated, consistent and complete information on your organisation's AI systems, it can be difficult to draw accessible insights and achieve a high-level overview of organisational AI risk, to make strategic decisions.

Agility

Without adaptive agility it is difficult to implement solutions to solve for current problems and flex to respond to enhanced requirements brought by increased regulation.

Without having a data driven strategy unpinning AI governance processes, it can be difficult to move swiftly and efficiently in response to requirements from different stakeholder groups, such as audits and investigations.

Drivers for good AI governance

There are multiple drivers for good AI governance, with clear benefits for prompt action. However, there is no set playbook. Each organisation must design its own framework to take account of its organisational strategy and risk appetite. There are early wins for organisations who act now, and who can be adaptive and agile to the evolving legal and regulatory environment.

Investing time now to take foundational steps to lay the groundwork to deal with more rigorous requirements as they emerge can save significant time and effort later. For example, raising awareness and taking stock of AI risk in current systems and learning from data, can help with more complex tasks of designing risk rating matrices across particular use cases and assist with developing strategies to mitigate AI risk. On the other hand, taking no action and waiting until future requirements become clearer, may mean that there is insufficient time to embed foundational pillars for good AI governance, with the risk that significant investment and management time is needed to catch up, along with the additional risk of financial penalties and reputational damage for getting it wrong.

Risks of non-action and benefits of prompt action

Risks of non-action

Material risk
Large fines up to 6% of global annual turnover or €30m (whichever is highest) resulting from EU's imminent proposed AI Act (the precise thresholds are still being agreed)

Reputational risk
Breaching public trust and confidence in your organisation arising from adverse AI system outcomes

Limiting growth
Missed opportunity to capture full benefits of AI within organisation

Benefits of early action

Capitalise on opportunity
Achieve confidence in adopting and realising the benefits of AI at scale, including improved operational efficiencies, reduced costs, and enhanced decision-making

Improve your bottom line
Improve public confidence and customer loyalty, increasing revenue growth and avoiding major costs

Improve AI project ROI (return on investment) by ensuring greater chance of successful implementation

Enhance your brand
As increasingly socially conscious customers and employees demand ethical practices from organisations, ensure you attract and retain customers and talent

Practical steps to take now

There are a number of practical steps that organisations can take now to assess their existing AI environment, identify any gaps in line with existing requirements and to plan their roadmap in readiness for future requirements.

Understand your current AI landscape

Assess your existing AI systems and baseline current maturity.

Review
Review existing governance arrangements for accountable roles and cross functional governance.

Define
Have a common understanding of AI systems and technologies in scope, bearing in mind that some simple automation techniques may carry highest risks.

Inventory
Create a consolidated inventory of identified AI systems and technologies, using a defined taxonomy.

Identify
Identify data points to capture to augment insights on what is driving AI risk.

Design
Design individual AI system risk assessment processes to ensure compliance with regulatory and policy requirements proportionate to your risk appetite.

Consider additional compliance assessments, such as bias evaluations and conformity assessments.

Assess
Assess individual AI systems and prioritise areas where there are gaps for specialist review.

Report
Develop reporting processes across your stakeholder groups.

Have a data driven approach

Have a data driven approach to accelerate delivery of risk management processes.

Capture data
Use digital questionnaires with structured data capture to consistently capture AI system information.

Automate workflow
Embed automated workflow to streamline processes to author, review and approve inputs and stage gates across multiple stakeholders and multiple teams – including referrals for external support.

Embed risk indicators
Tag responses to flag risk indicators that map to your risk appetite.

Consolidate data
Consolidate data capture to a centralised location for secure storage and access by multiple stakeholders.

Provide automated audit log to track status.

Visualise data
Design interactive dashboards to track KPIs and augment management reporting.

Have an easy view of aggregated or material AI risk across multiple data points (e.g., geography, AI system/ technology, business process impacted etc.).

Be agile
Be agile in designing and deploying solutions so you can solve for current problems.

Add new features as maturity grows, and as you learn from data and new requirements come online.

Benefits of acting now using technology and data to augment processes

Digital, centralised, AI inventory overcomes silos

- Break down silos with a central portal for all AI information
- Quickly obtain a full overview of all AI used in the organisation
- Accessed by all relevant stakeholders across departments to increase collaboration
- Be aware of who is responsible for each system

Efficient workflow enables data gathering

- Use digital questionnaires with inbuilt workflows for efficient, robust data gathering exercise across multiple departments
- Track status of questionnaire and approval process
- All answers logged for future reference and later analysis
- Access specialist advice with external referral mechanisms

Up to date audit trail to track compliance

- Information always up to date and all changes logged for audit trail
- Easily demonstrate compliance status to regulators
- Locate information quickly – data and technical documents associated to each AI use

Easily view and report on AI risk position

- Structured data enables advanced search, filter, and analysis of data
- Gain quick oversight or AI position or deep dive into individual use case
- Easily create reports for different stakeholders (audit, regulators, C-Suite)
- Quickly see gaps in compliance against defined fields

Grow maturity in AI risk position

- Baseline maturity and obtain insights of AI risk position based on data
- Identify gaps and plan roadmap to address key priorities
- Track risks and mitigations in place
- See progress of each system across its lifecycle

Focus on proposed EU AI Act

The proposed EU AI Act will have far reaching effects for many global organisations.

Organisations will need to understand what AI is caught by the regulation and implications for their AI governance processes.

The EU AI Act takes a risk-based approach. Organisations will need to classify their AI systems and practices and meet the compliance requirements that apply to each.

What AI is caught by the Regulation?

To fall within the current definition of AI System it must:

Have some autonomy

Infer how to achieve a given set of objectives

Influence the environment it interacts with, be it in a digital or physical dimension

The current definition:

"artificial intelligence system" (AI System) means a system that is designed to operate with elements of autonomy and that, based on machine and/or human provided data inputs, infers how to achieve a given set of objectives using machine learning and/or logic and knowledge based approaches, and produces system generated outputs (generative AI systems), predictions, recommendations, or decisions which influence the environments with which the AI System interacts" (Article 3(1)).

Classes of AI systems or practices

Art 5
Prohibited AI practices

Harmful subliminal behaviour distorting techniques

Harmful exploitative use against vulnerable, disabled or those in a "specific social/ economic situation"

Personality/social behaviour scoring causing detrimental treatment in another context or which is disproportionate to the behaviour

Non-targeted and disproportionate law enforcement use of real time biometrics in public spaces

Art 6
High risk AI systems

Annex II – listed products that already go through safety conformity assessment

Annex III – listed uses for biometric ID, educational/ employment evaluation, essential private/public service/benefits, law enforcement, migration, admin of justice/democracy

Compliance: stringent operational testing and deployment process and conformity assessment/ declaration

Art 4a
General Purpose AI systems

General purpose AI system which may be used as a component of high risk AI system

Compliance: most of requirements applicable to high risk AI, including conformity

Must cooperate with high risk AI providers to enable their compliance

Art 52
Declared use AI systems

Non-obvious human interactive systems

Biometric categorisation systems

Emotion recognition systems

Deep fakes

Compliance: providers/users must "inform of the operation of the system" to the decision subjects (crime detection exception)

Art 69
Lower risk AI systems

"AI systems other than high risk AI systems"

Compliance: Commission/ MS facilitate drawing up of voluntary codes of conduct for compliance with "one or more" of the requirements applicable to high risk AI systems tailored to specific uses

Other noteworthy modifications included in November 3, 2022, Council Compromise Text

- General purpose AI system obligations detail will follow through implementing acts (impact assessment required) within 18 months of AI Act coming into force.
- Social behaviour inappropriate scoring prohibition extended from just applying to public sector to also apply to private sector.
- AI systems specifically developed and put into service for the sole purpose of scientific research and development are excluded to ensure the Regulation does not affect scientific research and development activity on AI systems. Product oriented research activity by providers will also not attract the provisions of the Regulation.

Operationalising the Provider requirements: risk assessment and development timeline (High Risk AI Systems)

For AI systems classified as high risk organisations will need to operationalise provider requirements to ensure appropriate risks assessments and processes are in place.

1. Risk Management System

System purpose

System features

Checklist of regulatory requirements

Checklist of standards

Checklist of ethical risks

Library of mitigations

Understood by diverse participants/ stakeholders

Apply to system

2. Risk Assessment

Create initial blueprint/register of possible risks and constraints for project

Sensitive identified risks and how trade-offs calculated/justified

Iterative

Communicated to and signed off by senior management

3. Datasheets

Ensures data legality, sufficiency/ availability

Detailed human-generated metadata

4. Design Development Specification

Model selection

Design control

Harder for agile development

5. Training

Establishing design

Keeping results – feed into design and risk assessment

6. Validation

Focussed on the model

Keeping results – feed into design and risk assessment

7. Testing

Comprehensive

Keeping results – feed into design and risk assessment

8. Human Oversight

Creation of human oversight and intervention mechanisms

Sufficient explainability/transparency and limits

Override fall backs

9. User Instructions

Sufficient explainability/transparency/ residual risks

Creation of user instructions

10. Certify

Deploy

Finalise Art 11/Annex IV Technical Documentation – will have been built up through the process

11. Run

Capture of event logs – for when there is an issue

12. Feedback into Risk Assessment

Adjust as appropriate

13. Post-market Monitoring

Proactive monitoring/audit/testing as environment changes

Incident capture

Reporting – to authorities

14. Quality Management System

Documented system with accountability framework (risk management system, data/data governance, technical documentation, record-keeping, transparency, human oversight, accuracy, robustness, cybersecurity, and post-market monitoring)

Awards

Digital Legal Services award
CRR 194 Opinion Service (the Service)
FT Innovative Lawyers Awards Europe *2022*

Excellence Award
Best Use of Technology in a Law Firm
Canada Law Awards *2022*

Innovation of the Year (International Law Firm)
Global Legal Awards, Legal Week *2021*

Technology Venture of the Year
British Legal Technology Awards *2021*

Best use of Technology in a Law Firm
Canada Law Awards *2021*

Best of Legal 2021
Technology and data management
WirtschaftsWoche Best of Legal Awards *2021*

Winner
Most Innovative Use of Technology
The Lawyer Awards *2020*

Winner
Innovation and Technology Initiative of the Year
Legal Week, Asia Awards *2020*

Winner
Innovation of the Year (International Law Firm)
Global Legal Awards Legal Week *2020*

Gold Winner
Best Use of Technology in a Law Firm
Canadian Law Awards *2020*

FT Innovative Lawyers Report 2020
NRF Transform programme

Key contacts

Our multi-disciplinary team of leading specialists across technology and risk consulting, legal & regulatory and applied technology can help you to understand the risks and take advantage of the opportunities of AI within your organisation.

EMEA

Marcus Evans
EMEA Head of Information, Governance, Privacy and Cybersecurity
Tel +44 20 7444 3959
marcus.evans@nortonrosefulbright.com

Lara White
Partner
Tel +44 20 7444 5158
lara.white@nortonrosefulbright.com

Nadège Martin
Partner
Tel +33 1 56 59 53 74
nadege.martin@nortonrosefulbright.com

Jurriaan Jansen
Partner
Tel +31 20 462 9381
jurriaan.jansen@nortonrosefulbright.com

Christoph Ritzer
Partner
Tel +49 69 505096 241
christoph.ritzer@nortonrosefulbright.com

Michael Sinclair
Knowledge Director, Campaigns
Tel +44 20 7444 2344
michael.sinclair@nortonrosefulbright.com

Adam Sanitt
Head of Disputes Knowledge, Innovation and Business Support, EMEA
Tel +44 20 7444 2269
adam.sanitt@nortonrosefulbright.com

Peter McBurney
Co-Head of Technology Consulting
Tel +44 20 7444 5027
peter.mcburney@nortonrosefulbright.com

Madeline Bailey
Co-Head of Technology Consulting, EMEA
Tel +44 207 444 2806
madeline.bailey@nortonrosefulbright.com

Sarah Charig
Product Project Manager
Tel +44 20 7444 2010
sarah.charig@nortonrosefulbright.com

Chris Hendry
Senior Product Consultant
Tel +44 20 7444 3685
chris.hendry@nortonrosefulbright.com

APAC

Ross Phillipson
Senior Advisor, Risk Advisory
Tel +61 8 6212 3449
ross.phillipson@nortonrosefulbright.com

Stella Cramer
Global Co-Head of Technology
Tel +65 6309 5349
stella.cramer@nortonrosefulbright.com

Africa

Nerushka Bowan
Head of Technology and Innovation
Tel +27 11 685 8618
nerushka.bowan@nortonrosefulbright.com

North America

Daniel Farris
Partner-in-Charge
Tel +1 312 964 7730
daniel.farris@nortonrosefulbright.com

Maya Medeiros
Partner
Tel +1 604 641 4846
maya.medeiros@nortonrosefulbright.com

Anthony de Fazekas
Partner
Tel +1 416 216 2452
anthony.defazekas@nortonrosefulbright.com

Good AI governance is *essential* to mitigate against material AI risk.

NORTON ROSE FULBRIGHT

Norton Rose Fulbright is a global law firm. We provide the world's preeminent corporations and financial institutions with a full business law service. We have more than 3500 lawyers and other legal staff based in Europe, the United States, Canada, Latin America, Asia, Australia, Africa and the Middle East.

Law around the world

nortonrosefulbright.com

47375_EMEA – 11/22