

A Guide to Fundamental Rights Impact Assessments ("FRIA")

Under the EU Artificial Intelligence Act

December 2025

Funded by the European Artificial Intelligence & Society Fund

European Artificial Intelligence & Society Fund

This document is available under the Creative Commons license: [CC-BY SA 4.0 Attribution ShareAlike 4.0 International](https://creativecommons.org/licenses/by-sa/4.0/)

Acknowledgements

ECNL and DIHR would like to acknowledge the role of AlgorithmWatch and Michele Loi in the development of this guide, in particular the extensive contribution to the questionnaire template. Moreover, we received invaluable input from individuals and organisations who contributed their expertise, reflections and time on a voluntary basis, for which we are deeply thankful. We wish to extend our sincere thanks to: Gianclaudio Malgieri (eLaw, Leiden University), Isabelle Schipper (Netherlands Institute for Human Rights), Iris Muis and Julia Straatman (Utrecht Data School), Patricia Shaw (Beyond Reach Consulting Limited), and numerous civil society colleagues.

Note on methodology

This Guide is based on ECNL and DIHR materials (see [here](#)) and experiences, input from human rights experts and civil society organisations during 3 workshops in Berlin, Copenhagen, and Brussels, and public domain information such as academic literature (see, for example, [here](#) and [here](#)). The questions included in the Excel FRIA Template, specifically in the context analysis section, have been developed upon a review of similar questionnaires in AI impact assessment methodologies including the [Fundamental Rights and Algorithms Impact Assessment](#) used by the Government of Netherlands, the [Fundamental Rights Impact Assessment Model](#) of the Catalan Data Protection Authority, the FRIAct pre-deployment [questionnaire](#), and the Council of Europe [Methodology](#) for the risk and impact assessment of AI systems on human rights, democracy and rule of law. In the next phase, we invite organisations, deployers and developers to focus on piloting the Guide and we will continue updating it based on practical experiences. If your organisation is interested in implementing the Guide and/or have specific recommendations and suggestions, please contact Ioana Tuta (iotu@humanrights.dk) or Vanja Skoric (vanja@ecnل.org).

Table of Contents

1. INTRODUCTION	4
2. USING THIS GUIDE	5
A. A Holistic Reading of The AI Act	5
B. Drawing upon International Standards	6
C. Avoiding Common Pitfalls.....	6
3. CRITERIA FOR A MEANINGFUL FRIA	7
4. A STEP-BY-STEP GUIDE TO FRIA	9
Phase 1: Planning and Scoping.....	9
Phase 2: Assess and Mitigate Negative Impacts on Fundamental Rights.....	13
Phase 3: Deployment decision and public reporting	20
Phase 4: Monitoring and Review	22
Phase 5: Consulting Affected Groups and Stakeholders (cross-cutting).....	23
USEFUL RESOURCES	27
Impact assessment methods and practical guides	27
Academic literature.....	27



1. INTRODUCTION

The EU Artificial Intelligence (AI) Act, adopted in 2024, is an innovative legislative framework that regulates the development and deployment of AI through a risk-based architecture applying across industries and domains. With the protection of fundamental rights stated as one of its objectives, the AI Act requires organisations that deploy high-risk AI systems to carry out a fundamental rights impact assessment (FRIA) prior to its first use (see Box 1).

Beyond protecting fundamental rights, a well-designed and implemented FRIA carries many benefits: it can support organisations to develop and/or mature responsible AI governance frameworks, build trust and regular communication channels with external stakeholders, and minimise potential reputational and litigation costs. Moreover, in the case of public authorities, a FRIA creates incentives for the democratisation of decisions around AI uptake in critical areas where rights-holders interface with the State, e.g., education, healthcare, law enforcement.

The objective of this guide is to support organisations to conduct FRIAs in line with the EU AI Act and relevant international standards. As a decision-making milestone before the deployment of an AI system, a FRIA entails planning and scoping; structured deliberation on the severity and likelihood of negative impacts, including through the involvement of potentially affected people, their relevant proxies, and stakeholders with fundamental rights expertise; the implementation, and monitoring of mitigation measures; public transparency; and rigorous documentation throughout. It best delivers on its potential when anchored in an organisation-wide strategy for AI governance underpinned by dedicated policies and procedures, roles and responsibilities, resources and capacities.

BOX 1: THE ELEMENTS OF A FRIA ACCORDING TO [ARTICLE 27 OF THE AI ACT](#)

A FRIA should consist of:

- (a)** a description of the deployer's processes in which the high-risk AI system will be used in line with its intended purpose;
- (b)** a description of the period of time within which, and the frequency with which, each high-risk AI system is intended to be used;
- (c)** the categories of natural persons and groups likely to be affected by its use in the specific context;
- (d)** the specific risks of harm likely to have an impact on the categories of natural persons or groups of persons identified pursuant to point (c) of this paragraph, taking into account the information given by the provider
- (e)** a description of the implementation of human oversight measures, according to the instructions for use;
- (f)** the measures to be taken in the case of the materialisation of those risks, including the arrangements for internal governance and complaint mechanisms.

The main audience for this guide is deployers of high-risk AI systems under the AI Act, with the content specifically tailored for public authorities and bodies (see Box 2). Beyond this audience, the guide can be used by any organisation seeking to deploy AI responsibly and in accordance with fundamental rights.

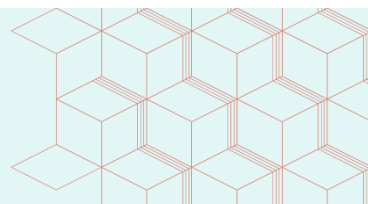
BOX 2: DEPLOYERS UNDER THE AI ACT

Three categories of deployers are required to conduct a FRIA:

- 1)** public authorities deploying high-risk AI systems in the areas listed in AI Act Annex 3, specifically biometrics, education and vocational training, employment, access to and enjoyment of essential private and public services, law enforcement, migration, asylum and border control management, administration of justice and democratic processes;
- 2)** private entities providing essential public services; and
- 3)** insurance and banking companies using AI to price life and health insurance and evaluate the creditworthiness of natural persons.

The document is structured as follows. Section 2 outlines the overall approach taken for the development of this guide. Section 3 outlines a set of criteria which can be used by deployers to assess whether the FRIA is grounded in a human rights-based approach. Section 4 includes a practical step-by-step guide to the planning and implementation of a FRIA. It should be read as a companion piece and explainer to the FRIA Template developed in a separate Excel workbook.

[Download the FRIA Template here.](#)



2. USING THIS GUIDE

a. A HOLISTIC READING OF THE AI ACT

The Guide places the FRIA obligation in the broader context of the AI Act by linking it with other relevant requirements on deployers of high-risk AI systems. For example, the requirements that deployers build internal AI literacy competencies (article 4) and put in place human oversight (article 26) have been highlighted as mitigation measures that can address negative impacts on fundamental rights. The requirement that deployers report serious incidents, including infringements on fundamental rights (article 26), has been connected with the expectation that deployers should establish a complaint mechanism.

Questions have emerged about the relationship between data protection impact assessments (DPIAs) required under the General Data Protection Regulation (GDPR) and FRIAs, specifically whether the two assessments should be integrated or kept separate when an organisation is required to comply with both. This guide has approached the DPIAs and FRIAs as stand-alone, separate, but complementary assessments (see Box 3).

BOX 3: THE RELATIONSHIP BETWEEN FRIA AND DPIA

Article 35 of the GDPR requires data processors to conduct a DPIA when the data processing poses a high risk to the rights and freedoms of persons. It has been pointed out that, because those organisations required to conduct a FRIA will also be required to conduct a DPIA under GDPR, the integration of the two assessments might become necessary. While there are obvious similarities between DPIAs and FRIAs, there are also some important differences. In practice, most DPIAs primarily address impacts on data protection whereas a FRIAs is meant to address impacts on all fundamental rights. Moreover, high-risk AI systems can have negative impacts that do not stem strictly from data processing, for example worker displacement due to new AI systems.

b. DRAWING UPON INTERNATIONAL STANDARDS

Endorsed by the UN Human Rights Council in 2011, the UN Guiding Principles on Business and Human Rights (UNGPs) have provided a normative foundation for integrating human rights in corporate risk management and impact assessment - a process known as 'human rights due diligence'. European policy makers have drawn upon the UNGPs to develop legally binding standards for responsible business conduct, such as the EU Corporate Sustainability Due Diligence Directive and Corporate Sustainability Reporting Directive. This guide draws upon the [UNGPs](#) in respect to, e.g., the methodology for assessing the severity of fundamental rights impacts, the central role afforded to stakeholder engagement, the effectiveness criteria for complaint mechanism.

c. AVOIDING COMMON PITFALLS

The Guide has also been informed by lessons learned emerging from the implementation of [algorithmic impact assessments](#) and human rights impact assessments. Typical concerns about impact assessments include: lack of adequate resource allocation and management buy-in, inadequate or pro forma stakeholder engagement, lack of organisational follow-up and monitoring of mitigation measures, insufficient public transparency, the timing of the assessment (i.e., after important operational decisions had already been taken). The legitimacy of FRIAs under the AI Act will depend on the extent to which such common pitfalls are overcome. For example, if an organisation conducts a FRIA as a desktop-based exercise without any external input and/or consultation, stakeholders might doubt the credibility of the findings. Or, if a FRIA is

implemented after a decision on the deployment of the AI system had already been taken, external stakeholders might be concerned that the organisation downplayed the severity of the potential negative impacts.

3. CRITERIA FOR A MEANINGFUL FRIA

CRITERIA	WHAT DOES IT MEAN IN PRACTICE?
Fundamental rights as benchmarks	<p>All impacts are assessed using fundamental rights standards as a benchmark. Relevant standards include the EU Charter on Fundamental Rights and European Convention on Human Rights, that can be complemented by specialised international human rights treaties such as the Convention on the Rights of the Child and Convention on the Rights of Persons with Disabilities.</p> <p>The focus of the assessment is on negative impacts on people (as opposed to positive impacts).</p> <p>The mitigation measures are consistent with, and not undermine, fundamental rights.</p> <p>The team conducting the FRIA draws upon fundamental rights expertise, e.g., input from fundamental rights organisations, review of relevant case law and authoritative recommendations from regional and international human rights bodies.</p>
Meaningful participation and non-discrimination	<p>Potentially affected rights-holders and/or their legitimate proxies are consulted in the identification and assessment of negative impacts and the design of mitigation and remedial measures (if necessary in the post-deployment phase).</p> <p>Organisations allow sufficient time to have meaningful exchanges with rights-holders and provide relevant and accessible information about the AI system and the objective of the consultation.</p> <p>Organisations consider and address possible barriers to participation, e.g., geographic distance, limited AI and digital literacy, lack of time and resources.</p>

	<p>Organisations take into account the differentiated needs of persons at heightened risk of harm and/or that benefit from protection under non-discrimination law such as children, persons with disabilities, migrants and asylum seekers, the elderly.</p>
Accountability	<p>Responsibility for the implementation, monitoring and follow-up of FRIA is assigned to particular individuals or functions within the company.</p> <p>Organisations maintain proper documentation and records to demonstrate the implementation of the FRIA.</p>
Access to remedy	<p>Impacted rightsholders have avenues whereby they can raise complaints about the AI system.</p> <p>Deployer-level complaints mechanisms are accessible and predictable and do not undermine existing judicial and non-judicial mechanisms.</p> <p>If a FRIA is updated after deployment, the organisation addresses negative impacts that have materialised and are in need for remediation.</p>
Transparency and access to information	<p>The FRIA summary required under the AI Act provides accurate, sufficient, and reliable information.</p> <p>As a matter of good practice, the FRIA findings are widely publicised and shared with the external stakeholders consulted during the process.</p>

4. A STEP-BY-STEP GUIDE TO FRIA

This part provides a structured approach for designing and implementing FRIAs as a companion to the FRIA Template which includes a practical set of questions that deployers can use to guide and document their deliberation and decision-making. The guide is structured in five distinct phases that address different aspects of a FRIA. While the first 4 phases can be followed in a sequential manner, the phase of stakeholder engagement has a cross-cutting relevance.



Phase 1: Planning and Scoping

Step 1.1. Consider the timing of the FRIA

Deployers are expected to conduct the FRIA prior to the first use of the high-risk AI system. For public authorities, public procurement provides a key junction to identify risks of negative human rights impacts and put mitigating measures in place to address such risks before harm occurs. For example, the FRIA can influence the [procurement](#) planning stage, including the supplier qualification and bid evaluation, as well as the contractual provisions with vendors. Conducting the FRIA only after an AI system has been purchased can undermine the deployer's ability to develop effective mitigation measures, specifically when these mitigation measures require implementation of technical measures by, and collaboration with, providers.

Step 1.2: Allocate a Realistic Budget to the FRIA

A FRIA involves cross-functional coordination, stakeholder engagement, iterative deliberation, mitigation measures, documentation, and monitoring. These activities require dedicated time and resources, necessitating an appropriate budget. The latter can be included in the broader cost-benefit analysis the organisation performs when considering the adoption of a new AI system.

Step 1.3: Establish the FRIA Team

The FRIA team should consist of experts from diverse disciplines, such as fundamental rights, engineering, AI safety and ethics, and internal functions, such as legal compliance, data protection, procurement, IT, operations. If the organisation does not have staff with fundamental rights competencies, it should assess how the FRIA team will be able to tap into that expertise (e.g, participation in training, consultation of external experts). While data protection expertise is relevant, the team should have knowledge of additional rights and jurisprudence relevant to the specific AI use case, such as anti-discrimination, due process rights, migration law, right to education, and right to health. Ideally, the team should be diverse, considering factors like gender and age.

The team's mandate is to advise the management on deploying the AI system in compliance with fundamental rights standards.

Assigning specific roles within the team, such as facilitating meetings, engaging stakeholders, and documenting findings, is recommended.

Three models for setting up the FRIA team can be envisioned:

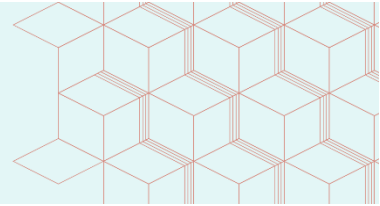
- **In-House:** A cross-functional team within the organisation conducts the FRIA.
- **Externalised:** A third party is commissioned to conduct the FRIA due to in-house capacity constraints or the need for independence. The deployer should ensure that the third party has the necessary expertise, and uses a fundamental rights-based methodology.
- **Hybrid:** An internal team works with a steering group of external experts for high-level advice and quality assurance. This model can be suitable for public entities deploying AI systems with large scale impacts.

The table below presents a list of pros and cons of these different models.

FRIA TEAM	PROS	CONS
IN-HOUSE	<ul style="list-style-type: none"> Ensures accountability through clear allocation of roles and responsibilities for implementation, follow-up and monitoring An opportunity for the deployer to build AI literacy skills and capacity on fundamental rights for staff Institutional memory on FRIA is retained, allows the generation of lessons learned and continuous improvement after deployment Better prepares the deployer to address issues that might appear during the monitoring of the AI system, e.g. serious incidents, responding to possible complaints 	<ul style="list-style-type: none"> If internal capacity/expertise is limited, FRIA can be implemented in a superficial manner External stakeholders might doubt the credibility of the assessment because of a perception of bias Internal constraints might prevent the team from thinking outside of the box
EXTERNALISED	<ul style="list-style-type: none"> Opportunity to tap into specialised expertise that might be lacking in-house Depending on the consultant chosen, it can mitigate against the perception of bias 	<ul style="list-style-type: none"> It can dilute the accountability of the deployer The assessment, and mitigation measures, might not be fully tailored to the deployer's institutional realities
HYBRID	<ul style="list-style-type: none"> Maintains institutional ownership and creates a structured approach for tapping into external expertise Corrects possible deployer bias through external input and feedback Bolsters the credibility of the assessment 	<ul style="list-style-type: none"> Operationalising the governance framework for the steering group and aligning timing and processes across internal and external stakeholders can be time and resource consuming

Step 1.4: Conduct a Context Analysis

[Download the FRIA Template here.](#)



The FRIA team's first task is to conduct a context analysis by going through a list of questions informed, amongst others, by the EU AI Act requirements. The context analysis questionnaire does not need to be filled out in one go – it might entail iterative discussions, identification of data gaps and need for further data collection, including via gathering input from internal and external stakeholders.

The questions are clustered in 3 categories:

1. **AI system deployment context:** Intended purpose of AI system; timeframe and frequency of use; decisions taken with the output of the AI system; applicable legal framework, including existence of judicial and non-judicial complaint mechanisms; identification of users and potentially affected people.
2. **AI system features:** Technical aspects; processing of personal data; available information about data quality, accuracy, and reliability.
3. **AI system governance:** Roles, responsibilities and information-sharing between the provider and the deployer; existence of internal policies and procedures on monitoring the AI system; measures for human oversight, AI literacy, and transparency towards affected people.

To answer these questions, the team can review:

- internal documents, e.g., pre-procurement analysis, cost-benefit assessments, data protection impact assessment, AI governance policies.
- documentation from suppliers/vendors, e.g., technical instructions, information about the provider's risk management systems and assessment of fundamental rights risks.
- publicly available information, e.g., concerns expressed by human rights organisations about similar systems, key fundamental rights stakeholders that can be consulted.

Some questions, for example those related to the AI system features, will be difficult to answer if the FRIA is conducted before the procurement of the system. Those questions are still useful to inform evaluation criteria to assess bids from potential suppliers. Deployers should prioritise suppliers that demonstrate a solid understanding of fundamental rights impacts and mitigation measures, and are committed to providing sufficient documentation and transparency.

The context analysis can prepare the ground for the identification and mitigation of negative impacts on fundamental rights in two interrelated ways. First, it can provide a first indication of possible sources of risks to fundamental rights, e.g., lack of consensus about the intended purpose of the AI system and how it will be integrated in the deployer's decision-making, low explainability of AI outputs, low capacity for meaningful human oversight, unclear/ambiguous regulatory framework. Second, it can indirectly suggest the types of mitigation measures that could be put into place, e.g., fundamental rights safeguards in contracts with suppliers, dedicated AI system monitoring processes.

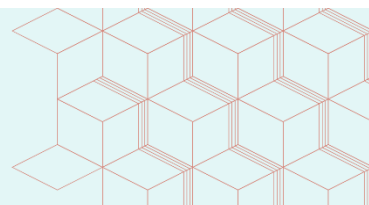
Relevant documentation that can support Phase 1:

- A workplan and timeline for the FRIA process
- Compilation of existing documentation
- Preliminary stakeholder mapping (see phase 5)
- If relevant, Terms of Reference (ToR) and/or contractual clauses for external consultants
- If relevant, ToR for the FRIA team



Phase 2: Assess and mitigate negative impacts on fundamental rights

[Download the FRIA Template here.](#)



In this phase, the FRIA team focuses on identifying the negative fundamental rights impacts that the AI system might have on people and deciding what mitigation measures are needed to address those negative impacts. The steps below correspond to the questions asked in the FRIA Template (in the 'impact assessment and mitigation' tab).

Step 2.1: Develop short scenarios and identify rights impacted

The FRIA team should develop brief scenarios outlining how the AI system might negatively impact people - including those affected by the decisions taken with AI outputs and those directly engaging with the AI system. The team could start by brainstorming all possible scenarios and select 3-5 for in-depth analysis, considering both typical and worst-case scenarios. Typical scenarios highlight negative outcomes under expected conditions, while worst-case scenarios focus on unexpected events or cumulative impacts that can exacerbate the severity of harm. The same scenario

could be analysed with a ‘typical’ and ‘worst’ case framing. For each scenario, all fundamental rights that could be negatively impacted should be identified. Each scenario is likely to implicate a few rights, as rights are often interconnected. For example, a biased AI system in education could impact the right to non-discrimination, the right to education, the rights of the child, and the right to privacy.

BOX 4: Illustrative typical and worst case scenarios

Scenario 1: Typical impact

During routine Mediterranean operations, a border protection agency deploys an AI detection system for people at sea across multiple zones. The algorithm consistently misclassifies fishing vessels operated by North African fishermen as "high-risk irregular migration" due to their vessel type, clothing patterns and movement behaviours that differ from European recreational boaters. Over three months, this generates 45% false positive alerts in certain zones, disproportionately triggering coast guard interceptions of legitimate fishing activities. Meanwhile, the system fails to detect a distress situation involving asylum seekers on a partially deflated raft during twilight hours with rough seas conditions underrepresented in training data. The delayed response results in deaths and severe hypothermia cases. The algorithm downgraded the alert because passengers' minimal movement (due to exhaustion) and lack of visible emergency signals were interpreted as "low risk." Coast guard officers, experiencing alert fatigue from frequent false positives and developing over-reliance on the system's risk assessments, fail to apply independent judgment in ambiguous cases.

Fundamental rights impacted:

- Right to life (Article 2, ECHR; Article 2, Charter): delayed rescue responses cause preventable deaths
- Prohibition of inhuman or degrading treatment (Article 3, ECHR; Article 4, Charter): exposure to life-threatening conditions at sea
- Bias that can potentially amount to indirect non-discrimination (Article 14, ECHR; Article 21, Charter): systematic bias against North African fishermen

Scenario 2: Worst-Case

Following a geopolitical crisis, migration routes shift dramatically as thousands flee conflict zones using unfamiliar departure points and vessel types absent from the AI's training data. The system, unable to recognise new patterns, systematically underestimates distress situations while simultaneously over-flagging routine maritime activities in affected regions. During this period, coast guard resources are misdirected to false alarms while genuine emergencies go undetected. The error affects operations across five Member States resulting in multiple mass casualty events. Simultaneously, the company supplying the system suffers a data breach, exposing biometric data, movement patterns and geolocation information of individuals to criminal networks and hostile state actors. Some Member States, facing domestic political pressure, begin using the system's "irregular migration risk" scores to justify blanket push-back policies without legal process.

Fundamental rights impacted:

- Right to life (Article 2, ECHR; Article 2, Charter): mass casualties from systematic detection failures and exposure of data to criminal networks
- Prohibition of torture and inhuman treatment (Article 3, ECHR; Article 4, Charter): prolonged exposure to life-threatening maritime conditions
- Right to asylum and non-refoulement (Article 18, Charter; Article 33, Refugee Convention): automated interdiction prevents access to protection
- Right to privacy and data protection (Article 8, ECHR; Articles 7-8, Charter): biometric data breach
- Right to an effective remedy (Article 13, ECHR; Article 47, Charter): opacity of automated decisions prevents accountability

To support the FRIA team with this step, FRIA Template includes a tab with non-exhaustive illustrations of how fundamental rights can be negatively impacted by AI systems ('Fundamental Rights' tab). It includes a list of all fundamental rights in the [EU Charter](#) and corresponding provisions in the [European Convention on Human Rights](#); a brief description of the content of each right; and a non-exhaustive list of negative impacts that might stem from the high-risk use cases subject to the FRIA obligation. The FRIA team should conduct additional research to identify recent fundamental rights developments relevant to the use case (see Box 5).

Box 5: Sources for research on AI and fundamental rights

- EU Fundamental Rights Agency [EU Charter of Fundamental Rights | European Union Agency for Fundamental Rights](#)
- EU Fundamental Rights Agency [Case Law Database | European Union Agency for Fundamental Rights](#)
- EU Fundamental Rights Agency [Data protection, privacy and new technologies | European Union Agency for Fundamental Rights](#)
- European Court of Human Rights, [Thematic factsheets](#) updated regularly and covering case law developments. Relevant themes include: prohibition of discrimination; freedom of expression; private life; surveillance at workplace; new technologies; right to a fair trial.
- European Data Protection Supervisor, [Artificial Intelligence | European Data Protection Supervisor](#)
- UN Human Rights Office, [Artificial Intelligence](#)
- Access Now, [Publications - Access Now](#)
- Amnesty International [Amnesty Tech - Amnesty International](#)
- Article 19 [Latest News - ARTICLE 19](#)
- European Digital Rights, [Publications - European Digital Rights \(EDRi\)](#)
- European Network of National Human Rights Institution, [Working Group on Artificial Intelligence](#)
- European Network for Equality Bodies, [Library | AI & Equality](#)
- European Center for Not-for-Profit Law [News & research | ECNL](#)
- Privacy International [Privacy International | Recent News](#)

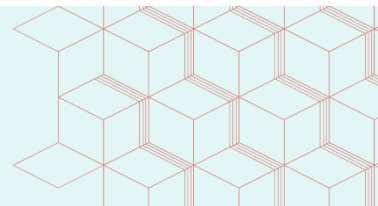


Fundamental rights are articulated in broad terms; it is through secondary legislation and case law that the precise scope of a legal obligation is clarified. The level of precision and certainty about what conduct reaches the threshold of violation varies significantly across fundamental rights. Moreover, in the case of violations

associated with the deployment of AI technology, there remains limited fundamental rights jurisprudence that practitioners can draw upon. In developing the scenarios, it is recommended then that the FRIA team take a holistic approach whereby all impacts that might undermine human dignity, freedoms, entitlements - irrespective of how neatly they can be correlated with legal obligations based on existing jurisprudence- are captured and discussed. That might mean, for example, considering impacts related to large-scale job losses as result of AI displacement of workers; de-skilling of educators and medical professionals as result of automation bias and over-reliance on AI system outputs; decline in the cognitive skills of children such as memory retention and erosion of critical thinking. Such a precautionary approach is warranted given the rapidly evolving nature of AI technology and the uncertainties about how it will interfere with the rights of individuals and groups in the short, medium, and long term. Rather than eschewing normative uncertainty, the latter should be seen as a feature of FRIAs that, if adequately documented and disclosed, can prompt broader further guidance and interpretations from fundamental rights bodies.

Step 2.2: Assess severity and likelihood of negative impacts

[Download the FRIA Template here.](#)



In the next step, the FRIA team assesses the severity and likelihood of each negative impact on fundamental rights.

Severity is evaluated based on several parameters:

- **Extent of interference:** Assess how the AI system interferes with a particular right, by prompting a reflection on the extent to which the deployer's conduct and negative impact might amount to an infringement of fundamental rights, e.g., based on secondary law (e.g., non-discrimination law, data protection law), case law, and non-binding opinions from regional and international human rights bodies.
- **Scope of impact:** Evaluate how widespread the impact is, the number of people affected, and whether persons in situation of vulnerability are impacted.
- **Gravity of impact:** Evaluate the seriousness of the harm (material, psychological, physical).

- **Irreversibility:** Evaluate how easy or difficult it is to restore those potentially affected to a situation at least the same as, or equivalent to, their situation before the negative impact.
- **Likelihood:** Consider factors affecting the likelihood of negative impacts, such as: the extent to which the provider of the high-risk AI system complies with the obligations in respect to risk management, data quality, accuracy, quality management system; the extent to which the deployer exercises meaningful human oversight; the degree to which the AI system will undergo further modifications; the deployers' interests, motivations and incentives; the extent to which malicious actors might seek to alter the AI system; changes in the legal framework.

The 'severity' and 'likelihood' questions in FRIA Template include closed-ended answers, with scales such as high, medium, low, that are used to calculate the prioritisation formula (see step 2.3 below). If the quantitative estimation used for prioritisation is not useful for the organisation or if there is a preference for a fully qualitative engagement, the pre-defined answers can be simply removed.

Stakeholder engagement to support and stress-test FRIA team's own assessment of severity and likelihood is critical (see phase 5). It should be noted that FRIA is a process of deliberation about what might happen in the future – as such, different people might have different perspectives and insights. Finding consensus is likely to require iterative conversations, identification of potential blind spots, and modifying underlying assumptions.

Box 6: Persons in situations of vulnerability

Persons in a situation of vulnerability are at a higher risk of interference with their fundamental rights and have fewer resources and capacities to seek remediation for such interferences and/or recover from associated harms, often due to societal or policy marginalisation. Individuals in vulnerable situations are likely to experience the same impact more severely. For example, an AI system setting health insurance premiums may disproportionately harm low-income individuals, leading to poorer health outcomes and higher poverty risks. Identifying the most disproportionately affected groups relevant to a specific deployment will be context-dependent. Examples mentioned in different sections of the AI Act include persons living in extreme poverty, ethnic or religious minorities (recital 29), persons under the age of 18 (art 9), people with disabilities (recital 132, art 60), historical patterns of discrimination, for example against women, certain age groups, persons with disabilities, or persons of certain racial or ethnic origins or sexual orientation (recital 65).

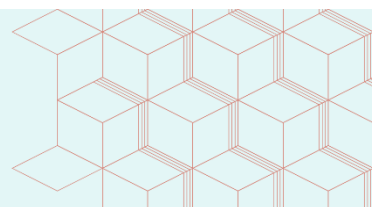
To identify persons in situations of vulnerability, deployers can pay attention to:

Personal characteristics: Grounds of potential discrimination such as sex, race, ethnicity, language, religion, political opinion, disability, age, or sexual orientation (Article 21, European Charter). While discrimination and vulnerability are not the same, vulnerability can often be caused or exacerbated by discrimination.

Power imbalance: Differences in status, knowledge, and socio-economic circumstances, e.g., between an irregular migrant and a border agency, or a welfare recipient and the authority deciding on benefits.

Step 2.3: Prioritise negative impacts

[Download the FRIA Template here.](#)



After answering questions about likelihood and severity, FRIA Template generates a prioritisation result (high/medium/low priority) and a risk matrix visualization. This allows for a comparative assessment of impacts on different rights based on their likelihood and severity. The prioritisation is underpinned by a formula embedded in FRIA Template and visible to users.

The prioritisation result helps deployers decide on the sequencing of prevention and mitigation efforts; it also ensures that deployers focus on those impacts likely to cause severe harm to people as opposed to those that might be most easy or convenient to mitigate. Having said that, while priority is given to the most severe, all impacts should eventually be addressed.

Step 2.4: Identify mitigation measures

For each identified impact, the FRIA team should determine existing or new measures to prevent or mitigate it. The AI Act suggests examples like human oversight arrangements and complaint handling procedures. Effective mitigation measures will evolve over time, especially at the technical level. Suggested classification of mitigation measures include:

- **Organisational measures**, e.g. meaningful human [oversight](#); complaint handling and redress procedures (see Box 7); public transparency; AI literacy; adopting a protocol response for dealing with requests for explainability.
- **Technical measures**, e.g., cybersecurity policies, managing and storing logs, ensure input data is relevant and representative.
- **[Contractual clauses for providers](#)**, e.g., requirements on input data quality and representativeness; requirements that providers cooperate with deployers in finalising and updating the FRIA; requirements for provision of documentation if deployers are faced with requests to explain the output of the AI system; requirements that providers inform deployers about any technical changes to the AI system that might impact its accuracy and performance.

Deployers should prioritise for implementation those measures that seek to avoid or prevent the negative impacts, followed by measures that reduce or minimise the negative impacts.

When developing the mitigation plan, the FRIA team should identify the internal functions and/or staff responsible for the implementation of the mitigation measures, and assess the need for additional capacities and resources, including for the monitoring of the effectiveness of measures.

Box 7: Complaint handling and redress procedures

Deployers should have in place procedures to receive and respond to queries, concerns and grievances, about the deployment of the AI system. Complaint mechanisms at the deployer level can complement judicial (e.g. courts) and non-judicial (e.g., equality bodies, data protection authorities, ombudsperson) remedial mechanisms by providing people with an accessible and speedy procedure to raise a broader set of concerns about the AI system. Under certain circumstances, such mechanisms can alert deployers to the existence of risks to fundamental rights or serious incidents which trigger notification obligations. Different considerations such as the nature of the deployer (e.g., education establishment, healthcare provider, law enforcement), pre-existing complaint mechanisms (e.g., whistleblower channels, worker complaint mechanisms), and the broader remedy context in the respective jurisdiction (e.g., availability and accessibility of non-judicial and judicial channels) can inform the design of such mechanisms. For example, upon mapping existing grievance mechanisms, deployers might decide that there is no need to build a bespoke mechanism but simply expand the scope of the existing one by ensuring it is available to a broader set of stakeholders, and that staff dealing with the complaints have sufficient AI literacy skills. While there might not be a one-size-fits-all approach across organisations, deployers can use the criteria below to assess the effectiveness of the complaint mechanism:

Accessibility

- The availability of the complaint mechanism is proactively communicated to potentially affected individuals.
- The complaint mechanism is open to a broad set of stakeholders, including civil society organisations, academics, and researchers.
- The complaint mechanism has a low threshold for complaint eligibility, without requiring proof of policy or law breaches.
- The interface chosen for the complaint mechanism (e.g., online forms, email, phone, in-person) considers barriers such as low digital literacy and limited internet connectivity.

Predictability and equitability

- The complaint mechanism has a public, clear process with a reasonable timeline for resolving complaints.
- The deployer provides access to additional information and expertise to complainants as needed.

Rights-compatibility

- Affected individuals using the complaint mechanism are informed of judicial and non-judicial remedies available in their jurisdiction for fundamental rights violations.
- If the deployer decides to provide a remedy, the remediation should be aligned with fundamental rights standards. Remedy can take different forms including restitution, compensation, rehabilitation, satisfaction, and guarantees of non-repetition.

Continuous learning and transparency

- The deployer uses the complaint mechanisms to learn and improve AI system governance.
- The deployer publishes information on the mechanism's performance, including the number and types of complaints received and their resolutions.

Relevant documentation that can support Phase 2:

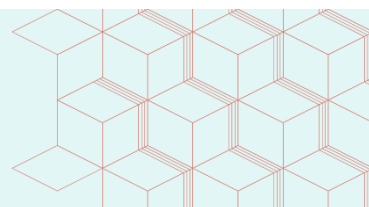
- Assessment of impact severity and likelihood as documented in the FRIA Template
- Documentation of consultation activities with affected people and other stakeholders, including objective of the consultation, people whom contributed, and how their input has been factored in the assessment of impacts and design of mitigation measures
- Mitigation plan with roles and responsibilities and expected budget
- Internal procedures for a complaint mechanism.



Phase 3: Deployment decision and public reporting

Step 3.1: Take a decision on deployment or non-deployment

[Download the FRIA Template here.](#)



After the assessment of negative impacts and identification of prevention and mitigation measures, the FRIA team should put forward a recommendation for or against the deployment of the AI system. A set of guiding questions to support that recommendation is included in the FRIA Template ('deployment tab'). Specifically, the FRIA team is asked to consider whether, after the implementation of prevention and mitigation measures, the impacts initially identified continue to be relevant, and/or whether their likelihood and severity of negative impacts has decreased. A decision on deployment should take into consideration whether the rights impacted are absolute or qualified. In the case of impact on qualified rights, deployers should assess the necessity and proportionality of any interference (see Box 8 below).

Box 8: Interferences with absolute and qualified rights

A few fundamental rights are considered to be absolute which means that no interference with these rights would be considered acceptable under any circumstances. Such rights include freedom of thought, the prohibition on torture, and freedom from forced labour.

Most of the rights, however, are qualified rights which means that interferences with such rights are acceptable under strict conditions. Examples of qualified rights include the right to privacy, freedom of expression, freedom of assembly.

Limitations on qualified rights are allowed if the law explicitly provides for the respective limitation; the limitation is meant to meet a legitimate objective of general interest such as public health, public security, public safety, the need to protect the rights of others; the limitation is necessary and proportionate regarding that objective. Such an assessment is highly context-specific, requires balancing different considerations and interests, and has been traditionally performed ex post facto by judges. It would be unrealistic to expect the FRIA team to conduct this exercise with the same level of precision and knowledge expected of judges – not least because of the inherent uncertainty of conducting such an assessment ex-ante in respect to potential interferences. This assessment, however, is useful to the extent that it can support the FRIA team in understanding the criteria against which the deployer's conduct would be assessed if confronted with an allegation of fundamental rights violation after deployment.

The following approach to deciding on deployment is recommended:

- In the case of interferences with absolute rights (even if severity and likelihood are assessed as low), deployment should not be recommended;
- In the case of interferences with qualified rights (when severity and/or likelihood are medium and/or high), the FRIA team should assess whether those interferences are necessary and proportionate and decide on deployment accordingly.

Step 3.2: Report publicly on FRIA findings

The AI Act requires deployers who are public entities to publish a summary of FRIA findings in the high-risk database. FRIA summaries for law enforcement and migration authorities will be published in a part of the database closed to the public (and only visible to market surveillance authorities).

The reporting obligation should be placed in a broader context whereby most EU countries have in place freedom of information laws guaranteeing the right to access information held by public authorities. A recent [decision](#) from the European Court on Human Rights clarified that the right to access to information includes a qualitative dimension whereby the information should be sincere, accurate, sufficient, and reliable. Against this backdrop, the FRIA summary should provide sufficient information in accessible/layman language to allow the broader public to understand:

- How the deployer is using the AI system and what types of decisions are being made with the outputs of the system;
- The FRIA methodology and limitations, including whether the deployer has engaged with affected people and/or their legitimate representatives and organisations with fundamental rights expertise;
- All identified negative impacts and the mitigation measures proposed, including justification of why the mitigation measures are deemed effective (e.g., allocation of resources, responsibilities, and capacities);
- Justification of the final deployment decisions;
- Existence of complaint mechanisms for raising concerns.

Deployers are encouraged to publish the FRIA summary on their own website and include additional documentation such as the FRIA team composition, the workplan, the types of data collected and analysed, the stakeholders engaged, the FRIA findings and deliberation, and the mitigation plans. Broadening transparency can build public trust in the quality of the process, and facilitate peer-learning on good practices.

Relevant documentation in Phase 3:

- Justification for the deployment decision based on the FRIA findings approved by senior management
- Documents for publishing FRIA findings.



Phase 4: Monitoring and Review

Step 4.1. Establish a monitoring process post-deployment

Following deployment, organisations should monitor the effectiveness of the mitigation measures and assess whether any changes and adjustments might be needed. A monitoring plan should provide an answer to the following questions:

- What, precisely, is to be monitored? For example, existence of and resolution of complaints about the outputs and predictions made with the AI system; the quality of human oversight; survey results on user satisfaction with the AI system.
- When / how long after the deployment of the AI system should monitoring activities occur? How often, at what intervals?

- Who, internally and/or externally, should conduct the monitoring activities? What is the governance structure, roles, procedures and organisation for the management of the monitoring plan? What is the process to agree on the findings of the monitoring and implement follow-up actions?

Step 4.2. Update the FRIA if initial circumstances have changed

The AI Act requires deployers to update the assessment if any of the required FRIA elements has changed or is no longer up to date. An update of FRIA can be considered in the following circumstances:

- a modification of the way in which the AI system is being used;
- the deployment has resulted in negative impacts on fundamental rights, including serious incidents, that have not been sufficiently or adequately captured in the initial assessment;
- the mitigation measures have not been working as expected and effectively;
- changes in legislation and new jurisprudence have changed the initial assessment of the legality of deployment.

The update of the FRIA would require the FRIA team to go through the steps outlined in phase 1-3 and 5.

Relevant documentation that can support Phase 4:

- Monitoring plan with information on what will be monitored, how frequently, roles and responsibilities, triggers for FRIA updates.



Phase 5: Consulting Affected Groups and Stakeholders (cross-cutting)

Existing resources such as the [Framework for Meaningful Engagement](#) (ECNL) and [Human Rights Impact Assessment Guidance](#) (DIHR) can be used to plan a meaningful and inclusive stakeholder engagement process.

Conducting FRIA requires more than desktop research and analysis: it demands meaningful engagement and consultation with the groups and communities whose rights may be affected by the AI system and/or their legitimate representatives. This is supported by Recital 96 of the AI Act which foresees that

deployers should involve relevant stakeholders in conducting impact assessments and designing mitigation measures. Moreover, a stand-alone right to participation is explicitly included in two human rights instruments mentioned in the AI Act in respect to people disproportionately harmed by AI systems. Specifically, the UN Convention on the Rights of the Child affirms children's rights to participate in, and influence, decision-making processes that can affect their lives (see [here](#) and [here](#)). Participation in all areas of life is a core principle underlying the [UN Convention on People with Disabilities](#).

In addition to potentially affected people, deployers can identify and engage with other stakeholders whose domain knowledge and expertise can support the FRIA deliberation by closing knowledge gaps and stress-testing assumptions about FRIA impacts. Such stakeholders can include: national human rights institutions, equality bodies, ombudspersons, data protection authorities, academics and researchers, civil society organisations and advocacy groups, industry bodies, professional associations, etc.

The consultation of affected groups and other stakeholders is relevant across all the FRIA phases as highlighted below.

Phase 1 Planning and scoping

The FRIA team identifies the stakeholders who should be engaged in the process (See box 9). Some preliminary interviews with stakeholders may also take place during this initial phase. Guiding questions to consider:

- Who is or might be affected by the deployment of the AI system?
 - Who will interact with the AI systems, e.g., police officers, healthcare staff, human resources staff?
 - Who are the individuals and groups affected by decisions taken with the AI system outputs, e.g., job applicants, persons applying for public welfare benefits, persons convening in public spaces, migrants, persons convicted of crimes?
 - Who are the individuals and groups that the AI systems will process personal data about?
 - Whom amongst these individuals and groups are most likely to be in a situation of vulnerability?
 - Do deployers have an obligation under the AI Act to inform these individuals that an AI system will be used to make decisions about them?
- What are the practical obstacles for directly engaging with affected people and how could they be overcome? Are there legitimate proxy organisations that represent the interests and rights of the affected people, e.g., trade unions,

migrants' rights organisations, children rights organisations, that could be engaged?

- In addition to affected people and their legitimate representatives, are there other stakeholders that should be prioritised for engagement?
- What are the time and resources needed to engage these individuals/groups?

Phase 2 Impact assessment and mitigation

Stakeholders should be meaningfully involved in the identification and assessment of impacts, as well as in the design of measures that effectively prevent, and mitigate impacts. Guiding questions to consider:

- Who should be responsible for organising the engagement with stakeholders? Is there a need for an external facilitator?
- What is the objective of the engagement, e.g., listen to the perspectives and concerns of potentially affected people about how the AI system might negatively impact them; get feedback on the mitigation measures; draw on specialised expertise to assess the acceptability of any remaining fundamental rights impacts?
- What would be the most appropriate in-person and/or digital modalities for engagement, e.g., public consultation, interviews, focus groups, surveys? Do these modalities factor in the needs of persons in situations of vulnerability, are the modalities age-appropriate and adapted to maturity of children?
- How to ensure the process is perceived as trustworthy as opposed to a check box exercise? What confidentiality protocols should be put in place?
- What type of information should be shared in advance, e.g. about the AI system, about the rights of affected people, the obligations of deployers, the negative impacts identified, to allow for a two-way exchange?

Phase 3 Deployment and public reporting

Consulted stakeholders should be informed about the FRIA results in a meaningful and accessible way. Guiding questions to consider:

- How should the consulted individuals and groups be informed about the FRIA findings, the decision on deployment, and how their input has been factored in decisions?
- How should the consultation process be documented and what information can be made public based on the confidentiality protocol agreed with the stakeholders?

Phase 4 Monitoring and review

Affected people should have mechanisms to raise concerns about the AI systems after deployment and be involved in the resolution of their concerns. Guiding questions to consider:

- Do deployers have effective complaint mechanisms through which affected people can raise concerns about the deployment of the AI system? Do those complaint mechanisms foresee for affected people to be involved in the design of remedial measures, if appropriate?
- Do deployers have adequate procedures in place to respond to requests for meaningful explanation of the role the AI system played in decision-making procedures?

Box 9: Example stakeholder mapping table

STAKEHOLDER CATEGORY TO ENGAGE WITH:
Primary stakeholders – affected people or groups:
<ul style="list-style-type: none">• Directly affected individuals and communities• Marginalised groups at heightened risk of harm, for example: children and persons with disabilities; racial, ethnic, and religious minorities; women and gender minorities; refugees, migrants, and people on the move• Legitimate representatives/proxies
Secondary stakeholders – intermediaries with potential knowledge about impact:
<ul style="list-style-type: none">• Different regulatory bodies• Civil society organisations and advocacy groups• Trade unions and professional associations• Academic researchers and human rights experts• Legal practitioners and judicial institutions• Human rights organisations such as equality bodies, national human rights institutions, regional/international organisations

USEFUL RESOURCES

Impact assessment methods and practical guides

- The Government of the Netherlands, Fundamental Rights and Algorithms Impact Assessment (2021), [Fundamental Rights and Algorithms Impact Assessment \(FRAIA\) | Report | Government.nl](#) and the Report on its piloting, [FRAIA in action | Report | Government.nl](#)
- Council of Europe methodology for AI impact assessment on human rights, democracy and rule of law (2024), [HUDERIA: New tool to assess the impact of AI systems on human rights - Portal](#)
- The Catalan Data Protection Authority, FRIA model: Guide and use cases (2025), [FRIA model: Guide and use cases](#)
- Danish Institute for Human Rights, Human Rights Impact Assessment of Digital Activities (2020), [Human rights impact assessment of digital activities | The Danish Institute for Human Rights](#)
- European Center for-Not-for Profit Law, Framework for Meaningful Engagement in AI Development (2025), [Overview](#) and [Framework for Meaningful Engagement 2.0 | ECNL](#)
- AlgorithmWatch, [Automated Decision-Making Systems in the Public Sector – An Impact Assessment Tool for Public Authorities](#) (2023)

Academic literature

- Karen Yeung, [Can risks to fundamental rights arising from AI systems be 'managed' alongside health and safety risks? Implementing Article 9 of the EU AI Act](#) (2025)
- Gianclaudio Malgieri and Cristiana Santos, [Assessing \(the Severity of\) Impacts on Fundamental Rights](#) forthcoming in Computer Law & Security Review (2024)
- Alessandro Mantelero, [The Fundamental Rights Impact Assessment \(FRIA\) in the AI Act: roots, legal obligations and key elements for a model template](#), in Computer Law & Security Review 54 (September 2024)
- Vanja Skoric, [Critical Criteria for AI Impact Assessment: An Aggregated View](#) (2023)

DANISH
INSTITUTE FOR
HUMAN RIGHTS



European Center for
Not-for-Profit Law

A Guide to Fundamental Rights Impact Assessments (“FRIA”)

Under the EU Artificial Intelligence Act

December 2025