Universiteit Utrecht

RIJKS ICT GILDE

FRAIA in action

Lessons learned from 15 FRAIA projects at Dutch government organisations

Utrecht University: Julia Straatman, Iris Muis
Rijks ICT Gilde: Giulia Bössenecker, Ruth Koole, Remco Draijer, Nout van Deijck, Isabel van Vledder
20 juni 2024

Management summary

The Fundamental Rights and Algorithms Impact Assessment (FRAIA)[1] is a tool to identify the impact of a specific algorithm on fundamental rights. Utrecht University developed the tool for the Ministry of the Interior and Kingdom Relations in 2021. The FRAIA can be used when a government organisation considers the development, procurement or use of an algorithm or to evaluate algorithms already in use.

In 2023, the Ministry of the Interior and Kingdom Relations instructed Utrecht University and Rijks ICT Gilde (the National ICT Guild) to conduct pilots with the FRAIA to gain experience and promote its use in government organisations. At the time of writing, the FRAIA is also receiving extensive political and international attention, including in the House of Representatives and in the context of the recently approved AI Act. The AI regulation includes an obligation to conduct a fundamental rights impact assessment for high-risk AI, which FRAIA may be able to fulfil on some level.

A total of fifteen pilots were conducted with various government organisations. The pilots included algorithms that had yet to be developed, algorithms under development and algorithms already in use, from *machine learning* to *business rules* algorithms. The organisations were distributed between Utrecht University and Rijks ICT Gilde and received guidance to complete the FRAIA in the context of their individual cases.

General findings

The pilots show that the FRAIA often leaves a positive impression despite prior scepticism. Participants appreciated the different legal, ethical and technical viewpoints and discussions that provided new insights. They consider the FRAIA document a useful tool for discussion about fundamental rights, data and ethics. However, participants found the process slow at times; certain questions were less relevant, and they struggled to engage all necessary functions.

Case identification

Some organisations that were eager to participate had difficulty identifying appropriate cases for the FRAIA, mainly because they could not always assess whether the technology involved an algorithm or affected fundamental rights. We propose the development of a pre-FRAIA quick scan to help determine if a case is

[1] See https://www.rijksoverheid.nl/documenten/rapporten/2021/02/25/impact-assessment-mensenrechten-en-algoritmes.

suitable for a full FRAIA process. This would help organisations work more efficiently and only conduct comprehensive assessments where necessary.

Intake and objective
A good intake and clear objectives proved crucial to a successful FRAIA programme. Organisations should consider in advance the purpose of the FRAIA, who the client is, who owns the document, and how the document will be used.

Answer requirements
There is a need for guidelines for answering questions in the FRAIA document because organisations do not always know when an answer is sufficient. It helps to establish the purpose of the completed document (e.g., justification or information) beforehand. There is also concern about the risk of seeing the FRAIA as a 'checkbox', where questions are answered as desired without due reflection. The FRAIA must not be used as a review tool or feel like a required element but as a means to facilitate ethical discussions regarding algorithms from which informed choices can be made.

Ethics and fundamental rights
The FRAIA focuses primarily on fundamental rights, creating the risk that ethical issues not directly related to fundamental rights are ignored. There is a need for strong moderation and flexibility within the process to discuss these ethical issues as they emerge.

Future view
After the pilot, organisations have varying expectations about using the FRAIA independently. Some plan to train staff internally, while others value external guidance. There is also a need for integration of the FRAIA with existing laws and regulations (e.g., the DPIA) and potentially an abridged version of the document for use as an accountability tool. Organisations are enthusiastic about the FRAIA, but further steps are needed to make the process more efficient and more widely applicable.

Recommendations
A complete overview of all recommendations can be found in Chapter 6. The recommendations cover the following topics:

1. *Before initiating the FRAIA*
   It is essential to determine on a case-by-case basis whether the FRAIA tool is necessary, especially for high-risk algorithms. Guidelines must be developed to determine at what point in the algorithm life cycle the FRAIA should be utilised. A key recommendation is the development of a pre-FRAIA scan. It is also important to clearly define the purpose, owner and version of the FRAIA document before initiating the process.

2. *Session structure*
   Guidance must ideally take place on-site, and the number of online participants must be limited. A good balance in the number of participants ensures effective discussions and results.

3. *Process supervision*
   A process supervisor is crucial for guiding the discussion and moving forward; the supervisor must be substantively independent to ensure objectivity. We recommend exploring the option of establishing a government-wide trainer pool for FRAIA process supervisors to support smaller government organisations.

4. *FRAIA form and content*
   Several aspects of the FRAIA, such as the job list and how the interactive PDF works, require revision based on the feedback collected. Integration of the FRAIA with other tools, such as the DPIA, can avoid duplication of effort and improve efficiency.

5. *After completing the FRAIA*
   The completed FRAIA document must be stored in a central location for easy reference by stakeholders. Treat the FRAIA as a dynamic document that can be modified and supplemented as the project progresses. Ensure proper version control and clarity on ownership and management of the document.

6. *Knowledge transfer*
   Consider how FRAIAs between different government organisations can be better streamlined, especially when municipalities use the same algorithms, for example. There is a strong need for a "FRAIA community" to promote collaboration and knowledge exchange. Expand the use of the FRAIA beyond the current pilot projects to disseminate knowledge and experience and increase support.

7. *Government, laws and regulations*
   Terms such as fundamental rights and human rights must be consistently used and interpreted in line with the AI Act. Explore whether and how to align the FRAIA with the requirements of a Fundamental Rights Impact Assessment (FRIA) under the AI Act.[2] Clarify the role of the FRAIA relative to other tools such as impact assessments, risk assessments, compliance assessments, and conformity assessments.

[2] See, for example, Article 27 of the AI Act: https://artificialintelligenceact.eu/article/27/.

Content

1. Introduction

1.1. FRAIA Introduction

The Fundamental Rights and Algorithms Impact Assessment (FRAIA)[3] is a tool to identify the impact of a specific algorithm on fundamental rights. It balances the expected positive impact of the algorithm against the expected negative impact on fundamental rights, after which an informed discussion can take place, and a decision can be made on whether or not to deploy the algorithm or whether modifications are necessary and desirable. The FRAIA enables an interdisciplinary dialogue conducted by the parties responsible for the development or deployment of an algorithmic system.

The FRAIA can be used when a government agency is considering the development, procurement, or use of an algorithm. The FRAIA can also serve as an evaluation tool for algorithms already deployed. The questions should ideally be discussed in a diverse team composed of a variety of specialisations and backgrounds. Utrecht University developed the FRAIA on behalf of the Ministry of the Interior and Kingdom Relations in 2020/2021. The authors are Prof. Janneke Gerards, Dr. Mirko Tobias Schäfer, Arthur Vankan and Iris Muis. Originally developed in Dutch as the Impact Assessment Mensenrechten en Algoritmes (IAMA), the English translation was published in 2022.[4]

The FRAIA is currently attracting extensive political interest. Motions have been filed and questions have been asked on this issue in both the Senate and the House of Representatives. A motion has been submitted in the House of Representatives to mandate human rights assessments (such as the FRAIA) prior to the use of algorithms when they are used to evaluate or make decisions about people (i.e., *high-impact* algorithms). [5] Another motion has requested publication of the results of these assessments in the Algorithm Register.[6] The Senate has requested that human rights assessments regarding the use of AI be conducted and repeated and that these tests be published.[7]

In response to motions and questions, the Cabinet, ahead of the AI Act, has stated that all new high-risk AI and algorithms to be developed by the government must be

[3] See https://www.rijksoverheid.nl/documenten/rapporten/2021/02/25/impact-assessment-mensenrechten-en-algoritmes.

[4] See https://www.government.nl/documents/reports/2021/07/31/impact-assessment-fundamental-rights-and-algorithms.

[5] See *Parliamentary Papers II*, 2021/22, 26643, no. 835 (via https://www.tweedekamer.nl/kamerstukken/moties/detail?id=2022D12329&did=2022D12329).

[6] *Parliamentary Papers II*, 2022/23, 36360VI, no. 14 (via https://www.tweedekamer.nl/kamerstukken/moties/detail?id=2023D25876&did=2023D25876).

[7] See https://www.eerstekamer.nl/toezegging/publicatie_mensenrechtentoetsen_ai.

subjected to a human rights assessment[8] and considers it desirable that this assessment be repeated for the use of AI. The results of the human rights assessment can be published in the Algorithm Register. The Cabinet states that a human rights assessment will be included in the Algorithm Framework as a measure to help determine in a timely manner whether laws and regulations have been complied with and whether systems align with public values. The Algorithm Framework provides a handy overview of key standards and suggested tools to help governments more easily comply with the standards.[9]

The AI Act adopted in March 2024 is important in this context. This European legislation requires organisations to conduct a Fundamental Rights Impact Assessment (FRIA) when an AI system is classified as high-risk. It is relevant to compare the requirements of the AI Act with the information in the FRAIA. The FRAIA could fully or partially satisfy such obligations in the future.

1.2. Reason

The Ministry of the Interior and Kingdom Relations (hereinafter: Interior Ministry) commissioned Rijks ICT Gilde: (hereinafter: RIG) and Utrecht University (hereafter: UU) to supervise a number of FRAIA processes with government organisations. These were implemented in 2023, whereby eight processes were supervised by UU and eight by RIG. The participating government organisations volunteered for the pilots.

The project aimed to promote the use of the FRAIA and gain experience with its deployment through the pilots. This is also reflected in the Values-Driven Digitalisation work agenda (line 2.1: Safeguard Public Values and line 3.3: Regulate Algorithms).[10] These pilots build knowledge and expertise on issues regarding algorithms and human rights in general and the use of the FRAIA in particular. The idea is that Dutch government organisations will be sufficiently prepared when the AI Act comes into force and mandates the FRAIA or a similar assessment.

1.3. Goal

The purpose of the assignment is twofold:

- **Encourage the use of the FRAIA among government organisations (particularly municipalities).**

[8] Parliamentary Papers II, 2022-23, 26 643, no. 1056.

[9] See https://open.overheid.nl/documenten/27ada380-da3e-47ce-8c67-61044d250ee6/file.

[10] Appendix to *Parliamentary Papers II*, 2022/23, 26642, no. 940 (see: https://www.tweedekamer.nl/kamerstukken/brieven_regering/detail?id=2022Z21101&did=2022D45419).

Offering the option for external guidance encourages government organisations—particularly municipalities—to start using the FRAIA and gain experience with the tool. This removes barriers that some organisations may experience, such as concerns that they lack the in-house knowledge required to use the FRAIA. The idea is that after completing the process with external guidance, organisations can continue using the FRAIA independently.

- **Collecting user feedback on the FRAIA.**

Collecting user feedback also results in *best practices* and recommendations that can be shared widely. The feedback is important for improving the FRAIA and should also be considered in light of new obligations, such as the AI Act with which the FRAIA can be aligned. The requested feedback related to:

  - Overall impression
  - Concrete areas for improvement
  - Future vision
  - Implementation (substantive and procedural)

1.4. Reading guide

Several authors from two organisations have compiled this report: Utrecht University (UU) and Rijks ICT Gilde (RIG). The two parties facilitated FRAIA sessions independently. As such, the *Method* section (see Appendix) is divided into two parts. The other chapters were written collaboratively. For the sake of readability, we opted to consolidate the results of the two parties in a single chapter (*Findings*).

1.4.1. Terminology

To avoid confusion, it is important to briefly explain certain terminology.

1.4.1.1. Algorithm and AI

The terms *algorithm* and *AI* are frequently used interchangeably. An algorithm is defined as a set of rules and instructions that a computer follows automatically when running calculations to solve a problem or answer a question.[11] An AI system is a machine-based system that, for explicit or implicit objectives, infers, based on the inputs it receives, how to generate output, such as predictions, *content*, recommendations or decisions that can influence physical or virtual environments.

[11] See https://algoritmes.overheid.nl/en.

Different AI systems vary in their levels of autonomy and adaptiveness after deployment.[12]

As previously mentioned, the AI Act contains a legal requirement to conduct a Fundamental Rights Impact Assessment (FRIA) when high-risk AI is involved. This legal requirement does not apply to all algorithms because not all algorithms contain AI. The FRAIA is suitable for both algorithms and AI. The potential requirement to conduct a FRAIA for impactful algorithms in the Netherlands would surpass the requirements of the AI Act.

1.4.1.2. FRAIA session and FRAIA section

A *FRAIA session* refers to a session at the participating organisation during which part of the FRAIA is discussed. The content of each session may vary for each organisation. A *FRAIA section* refers to a specific part of the FRAIA (the FRAIA is divided into four parts). For example, session 2 of the FRAIA process at Municipality X does not necessarily mean that only FRAIA section 2 was discussed. The breakdown of the sessions by each participating organisation is described in the Method (see Appendix).

1.4.1.3. Human rights and fundamental rights

The FRAIA only examines the risk of affecting *fundamental rights*. Accordingly, this report uses the term *fundamental rights*. [13]Fundamental rights are the rights enshrined in the constitution of a specific country, in contrast to human rights, which apply to all people worldwide.[14]

1.4.1.4. Human rights impact assessment and human rights check

A *human rights impact assessment* is conducted to identify and assess any intended and unintended impact on human rights. A *human rights check* determines whether a system or application meets the legal and ethical human rights standards. The FRAIA is a fundamental rights assessment. As previously stated, the AI Act uses the term *fundamental rights* with respect to assessments. The aforementioned motions use the terms human rights check and human rights assessment. These two very similar terms are not always interchangeable.

[12] See https://oecd.ai/en/wonk/ai-system-definition-update.

[13] See https://www.rijksoverheid.nl/onderwerpen/mensenrechten/mensenrechten-nederland.

[14] See https://www.rijksoverheid.nl/onderwerpen/grondwet-en-statuut/grondwet.

2. Findings

This chapter presents UU and RIG's findings with a breakdown per theme. Feedback from participating organisations was abstracted and made untraceable to the respective organisation.

A rough distinction can be made between substantive and procedural feedback in the findings. The rest of this chapter will explain this distinction. Some of the subtopics include anonymous illustrative quotes from the participating parties.

To start, it is worth naming the participating government parties. These were as follows, in alphabetical order:

Government organisation	Guidance from RIG/UU	Domain
Tax and Customs Administration	RIG	International services
Municipality of Almere	RIG	Social Domain
Municipality of Assen	UU	Community approach
Municipality of The Hague 1	RIG	Social Domain
Municipality of The Hague 2	RIG	Logistics and transportation
Municipality of Enschede	UU	Social Domain
Municipality of Haarlem	UU	Enforcement and Public Order & Safety
Municipality of Den Bosch	UU	Enforcement and Public Order & Safety
Municipality of Veenendaal	UU	Social Domain
Royal Netherlands Marechaussee	RIG	Security
Logius	UU	Electronic identification
Education Inspectorate	UU	Supervision
Open State Foundation	RIG	Generative AI and the public sector
Province of North Brabant	RIG	Ecology and environmental protection
Province of South Holland	UU	Spatial Domain

Table 1 Overview of participating government organisations. The Open State Foundation's process has not yet been fully completed at the time of writing.

Initially, only municipalities were invited to participate in the pilots. However, the Interior Ministry, UU and RIG received too few applications from municipalities, so the

scope was broadened to include other government organisations. Even after broadening the scope, it still proved challenging to find participants for all the processes.

This may have been due to government organisations not being very aware yet of or feeling little involvement with algorithm evaluation or that municipalities did not yet have a suitable case to complete the FRAIA with. When a suitable case was found, there was a lack of internal support for reviewing the algorithm with the FRAIA. The FRAIA is still unknown to many government organisations, so the sense of urgency for the pilot was lacking among some of the organisations approached.

Finally, several government organisations were eager to participate but could not make time for it until 2024 due to staff shortages.

2.1. Substantive feedback

2.1.1. General feedback

The general impression of the participating organisations can be summarised as follows:

- Participants who were sceptical about the usefulness of the FRAIA or the sessions were *always* positively surprised afterwards. For example, they stated that they appreciated the different perspectives or that the in-depth discussions yielded different insights than expected.

- Participants expressed interest in hearing how other roles made certain decisions during the algorithm development process.

- Participants claimed to better understand the importance of discussing fundamental rights and ethics afterwards.

- Participants find the FRAIA document very suitable as a tool to facilitate conversation.

- Participants found the FRAIA process time-consuming.

- Participants do not find all questions equally relevant and would like to see an abridged version of the FRAIA.

- Participants often found it difficult to schedule a time for the process with all participating roles.

- Most organisations feel that some questions are not clear enough, creating a risk of different interpretations of the questions.

- The FRAIA cites the concept of *public values* several times. The difference between public values and fundamental rights was not initially clear in many cases.

These frequently mentioned feedback points are further explained in the following themes and can be illustrated by the following quotes:

"FRAIA has brought things to light that an organisation needs to take action on, even in other areas (beyond this project)."

"The pilot set up with interview sessions is very valuable. To merely put this in writing does not seem to do justice to the conversation and dialogue organisations want to initiate internally and externally."

2.1.2. Identifying an appropriate case study

The pilots have shown that organisations regularly struggle to identify an appropriate case for the FRAIA. This can be broken down into several aspects:

- *Algorithm.* Organisations are sometimes at a loss as to whether a particular technology concerns an algorithm. It is not always clear to such organisations which technologies should and should not be subjected to the FRAIA.

- *Fundamental rights and risk.* The FRAIA lends itself best to algorithms that may impact fundamental rights, known as *high-risk algorithms*. Some organisations indicate that it is not always clear in advance whether an algorithm actually affects a fundamental right. Although the FRAIA also serves partly to determine this, organisations want to avoid completing a full FRAIA process for each algorithm (high and low-risk) before knowing whether their algorithm affects fundamental rights. This touches on the broader discussion around the definition of high-risk algorithms and whether there should be a quick scan for the FRAIA. This quick scan or "pre-FRAIA" should clarify whether a case is suitable for the FRAIA or whether the FRAIA is unnecessary. The need for a quick scan is apparent among all participating parties. However, participants also stated that, while the FRAIA process is valuable, it is not tenable to complete a full FRAIA process for every potentially high-risk algorithm in the future.

- *AI Act.* The recently adopted AI Act includes a requirement to conduct a Fundamental Rights Impact Assessment (FRIA) for high-risk **AI systems**. However, the FRAIA can be run on a wide range of algorithm types. These may involve AI, but that is not necessary. Several organisations have provided feedback that they are under the assumption that the FRAIA can only be conducted on high-risk AI systems as defined in the AI Act. This is not the case.

An example of how participating organisations will address this is as follows:

"We will start by running the FRAIA on all our algorithms to gain experience. After that, we can start distinguishing between different types of algorithms based on complexity, for example. We could work with other municipalities to create something like a decision tree."

"The Privacy Officer may want to start incorporating preliminary questions about the FRAIA into the existing pre-scan for the DPIA and BBN (information security)."

2.1.3. FRAIA intake and objectives

In general, a proper intake is important for the successful implementation of a FRAIA. The FRAIA can be conducted on projects that are at different stages. It is not only important to identify an appropriate case but also to discuss the desired objective. For example, the outcomes of a FRAIA will differ for a start-up project in which important choices still need to be made compared to a project approaching implementation. Whereas the FRAIA can mainly provide direction on how to develop an algorithm for a start-up project, it can play an important role in a GO/NO GO decision for a project in its final phase.

2.1.4. Purpose identification

Upon conclusion of the pilots, it turned out that not all organisations had thought about the purpose of conducting a FRAIA on the case in question prior to the sessions. This caused the pilot to be more of an introduction to the FRAIA process for some organisations. This is appropriate for the pilots conducted, but without actual targeting (in situations outside the pilot), certain essential questions remain unanswered and are not asked in the FRAIA. Questions that may be raised during purpose identification:

- Who orders a FRAIA, and for what purpose?

- Who receives the completed FRAIA, and what happens to it?
- Who is the document owner?
- Is it a one-off task, or should the document be kept up to date? If it should be kept up to date, who is responsible for that?
- Will the FRAIA be used as input for a decision (e.g. Go/No go)? If so, which decision and by whom?

This was also emphasised by one of the participating parties:

> *"Purpose identification is essential. The clearer you describe it, the easier it is to complete the rest of the questions. Describing the purpose clearly is not easy. We thought about it for quite a long time to formulate it correctly. The importance of the purpose identification could be emphasised even more."*

2.1.5. Answer requirements

Several organisations stated that they could not accurately assess when a question had been adequately answered. The FRAIA contains no guidelines for this. It is important to determine the ultimate purpose of the completed FRAIA document. If the document will be shared with citizens, for example, consideration should be given to potential jargon and confidential information. If the document is used to justify the decision process behind the algorithm in question, the questions will likely need to be answered more comprehensively (i.e., with more explanation and background information) than if the document will be used only by those directly involved. Organisations expressed a need for guidelines on this topic.

In addition, some organisations are not clear about the status of their responses. The FRAIA is not an assessment tool, but rather a prerequisite for some organisations to deploy certain algorithms. This could result in the FRAIA being treated as a 'checkbox', where questions are simply answered as required. The FRAIA does not test; it facilitates. If the outcome of the FRAIA is that an algorithm has a disproportionate impact on human rights or that an algorithm is unnecessary or subsidiary, there is reason to revise or cease the development or use of that algorithm. This does not apply vice versa: a positive conclusion based on the FRAIA does not automatically mean that the development or use of the algorithm should proceed. As such, the

FRAIA is insufficient as a tool for accountability/assurance/conformance. This distinction is not currently clear enough.

2.1.6. Ethics and fundamental rights

In some cases among the FRAIA pilots, fundamental rights were not affected but other problem areas unrelated to fundamental rights surfaced. In some cases—despite a lack of infringement on fundamental rights—such ethically objectionable sticking points surfaced that the case had to be reviewed with the responsible parties. The FRAIA does not automatically catch such issues, given that it is ultimately a *fundamental rights* test. If a case does not infringe on a fundamental right, there may be a risk of other non-fundamental rights-related considerations being disregarded.

Using this example, it is important to emphasise that the FRAIA is a comprehensive fundamental rights assessment. While it may capture many ethical sticking points, it is worth noting that a fundamental rights test will not automatically be adequate for all possible ethical issues that may arise around an algorithmic system. This can be mitigated through strong moderation and the flexibility to raise these issues as they arise.

Some organisations disclosed that it is sometimes difficult to determine in advance whether an algorithm will affect fundamental rights. One organisation said the following:

> *"The section on human rights, Section 4, should be completed first if it is unclear whether fundamental rights will be affected. This enables an immediate decision as to whether the rest should be filled in at all."*

However, the idea behind the design of the FRAIA is that all parts are interrelated, and conclusions cannot be drawn solely from completing Section 4.

2.2. Procedural feedback

2.2.1. Leaders

Before sharing further findings, it is worth briefly reflecting on the roles of the process initiators. These roles are shown in the table below.

FRAIA process initiator role

CIO office advisor
Data ethics consultant
Ethics business analyst
CDO
CIO
Data protection officer
Privacy advisor
Strategic information manager
Science officer

Table 2 FRAIA processes: roles of the initiators within the organisation.

This refers to the individuals (and their corresponding positions) who have contacted UU or RIG, either through the Interior Ministry or through the calls from UU and RIG themselves. These individuals initiated the FRAIA pilot within their respective organisations.

2.2.2. Roles and attendees during the FRAIA process

The FRAIA processes were always completed with interdisciplinary teams, as the tool explicitly recommends. All the roles involved in the FRAIA processes are listed in the Appendix. The various sections cover legal, ethical and technical aspects. The job function table on the second page of the FRAIA serves as a guide. Many questions were asked about this list during the intake with organisations:

- Organisations were usually under the assumption that all the roles listed should be present.
- Organisations had questions about the added value of certain roles, such as an HR officer or a citizens' panel.
- Organisations suggested that some roles were missing, such as a data protection officer for all sessions.
- Not all organisations had all the roles.

The questions and comments above suggest that the list may cause confusion and that organisations need more substantiation or accompanying explanations.

When an organisation was unable to independently determine which roles to invite, it was advised to invite those employees able to answer the questions in the FRAIA for the particular case. Experience shows that this naturally results in an interdisciplinary team.

2.2.3. External developer

Algorithms within government organisations are not always developed internally. It may be important for government organisations to invite external developers when going through a FRAIA process. The FRAIA pilots have also shown the value of involving the external developer, particularly in answering the technical questions about how the algorithm works in the FRAIA document.

At the conclusion of the sessions, all external developers generally stated that the sessions were instructive and that they understood the importance of the FRAIA for high-risk algorithms. Again, it was emphasised that knowledge of the FRAIA is invaluable when frequently developing for government organisations. After all, the algorithms by these developers will have to meet government requirements and possibly also the FRAIA.

However, several external developers stated that the joint meetings and further completion of the FRAIA document were very time-consuming, and the additional time must be charged to the client:

> *"It is a very time-consuming task where, in a preliminary phase, many issues may not be sufficiently clear ahead of time to answer properly, and there is a risk of a catch-22."*

Lastly, organisations stated that participation by the external developer provided a better understanding of how the algorithm in question works. This improved understanding may enable organisations to better identify algorithm opportunities and risks.

2.2.4. Group size and roles

Group size is an important consideration in the successful completion of a FRAIA process. Too few participants may result in a knowledge deficit, leaving some FRAIA questions unanswered. However, if there are too many participants, the excessive

number of people can cause excessive discussion, underhandedness, and potential redundant attendance due to duplicate roles, for example. The FRAIA pilots saw a wide range of group sizes, ranging from three to twenty participants per session.

For some organisations, not all affiliated roles were clear on their added value during the sessions. They contributed to the group size but not much to the discussion. It is important to clarify in advance what each role will contribute to and gain from the sessions. An excess of superfluous functions leads to a messy overall picture, where a small group leads the discussion, and the rest of the participants are merely present. Superfluous functions lacked knowledge about the case at hand, leading to frequent, time-consuming discussions about semantics. One organisation said the following:

> *"Some participating functions had a knowledge gap and delayed the process. This led to frustration among the well-informed participating roles."*

> *"I have been told to attend these sessions, but I don't really understand why."*

At the same time, some organisations only had the essential functions (for answering the questions) attend the sessions. This resulted in limited discussion on certain topics, and the organisation was unable to ask critical questions about those topics.

The authors found it difficult to make a statement about the ideal group size, as they disagreed on that matter. However, the final group size will vary by FRAIA process and will be case-dependent. This makes prescribing guidelines for this variable difficult, and some organisations would have liked more certainty.

2.2.5. The average duration of a FRAIA session

The starting point for conducting the FRAIA within the pilot project was to stick to the standard format of five hours in total.[15] This format was deviated from in some cases during the pilots. In several cases, the five hours proved too ambitious to fully complete the FRAIA (i.e., fully capture the impact of the examined algorithm); the thickness of the FRAIA document and the time frame for completion proved an obstacle for some organisations. This does not mean that the FRAIA sessions have

[15] The five-hour duration covers completing the FRAIA itself. The intake interview (30 minutes) and the evaluation session (30 minutes) are not considered here, nor are other peripheral matters such as scheduling.

not been effective. In many cases, the sessions led to increased awareness and an initial assessment of risks. Where there was insufficient time for a detailed discussion, actions were formulated. The five hours were ample for many cases, and some even had time left.

Whether the suggested five hours is sufficient depends mainly on the predetermined objective. If the FRAIA is intended to raise awareness around the risks of a specific algorithm, five hours is sufficient. This assumes that a lot of the necessary work has been completed beforehand. When choices still needed to be made or when there was a lot of discussion about a choice made, participants felt that the proposed timetable was sometimes unfeasible. There is often no right or wrong when discussing ethical issues, which can make these discussions lengthy. It is worth noting that these discussions should be seen as part of the project and not directly attributable to the FRAIA. The FRAIA merely facilitates the identification and discussion of the ethical issue in a broad forum.

RIG also noted that it was sometimes desirable to delve deeper into a particular subject but that this often proved impossible due to time constraints. The questions in the FRAIA often remain relatively superficial, and it can be quite a challenge to formulate a complete and clear answer to them. In some cases, further investigation into data quality or algorithm performance may be necessary, for example. During the sessions, there was often only time to talk about these topics abstractly, whereas a deeper discussion in relation to the algorithm can provide more useful insights and discussions. At the same time, UU views this as a task for the respective organisation, and these five hours can serve as an impetus for further discussion or actions.

The necessary time investment also depends on the following factors: the complexity of the case, having the right knowledge (and the right number of people) present, and having good and efficient guidance and moderation. These points will be discussed further later. A specific explanation of the moderator's role can be found in the Appendix.

2.2.6. Multi-day and and single-day sessions

The five-hour time investment for completing the FRAIA document can be categorised in several ways. During the pilots, sessions were most often spread over three days and otherwise over two days. On a few occasions, all sessions were scheduled on one day. However, it is worth noting that experience (outside the pilots) has shown that feelings can sometimes run high during a FRAIA process. If all sessions are done in one day, there is no time to cool off or discuss issues outside of the process. There is also no time for an interim assessment of whether other participants need to be there after all, nor can interim evaluations be conducted for other reasons.

We saw different versions of FRAIA processes when sessions were split out over multiple days. In some processes, the sessions were far apart (more than a month). In others, the sessions took place very close together (all sessions in one week). Again, striking a balance is key; too much time between sessions does not benefit the process. Participants may have forgotten some of the material discussed in the previous session, and time that could have been spent on the current session will need to be spent summarising. On the other hand, if the sessions are too close together, there is a risk that there is not enough time for reflection and interim evaluation. We recommend leaving one to two weeks between sessions, keeping in mind practical concerns such as scheduling, if possible.

Some organisations had the following to say regarding staggering the sessions:

"This setting [doing everything in one day] was perfect for our issue. However, I think it would be difficult to determine whether the same could be done for other cases."

"When it is a more complex matter, with more participants, more interests or a more complicated product, it becomes more difficult. You need more time to hear everyone. It probably wouldn't be possible in one day, in that case."

2.2.7. *In-person, online and hybrid sessions*

Most sessions took place in person. On a few occasions, participants were unable to attend in person and signed in online. Most of the participants still attended in person.

During the hybrid sessions, the online participants struggled to fully participate in the discussions. Discussions largely take place between participants in the room, making it difficult to break in from a distance. As such, in hybrid sessions, it is important for the group and the moderator to actively engage with online participants.

UU and RIG prefer all participants to be physically present to facilitate a smooth conversation about the FRAIA questions and not miss any input.

2.2.8. *Future vision*

After completing the FRAIA process, the participating organisation has been thoroughly introduced to the FRAIA. That leaves the question as to whether this tool can now be used independently and what the next steps will be. Whether the FRAIA

can be used independently after the pilot ends varies per organisation. Several organisations have suggested training people in-house to independently apply the FRAIA within the organisation. Other organisations are interested in UU's FRAIA train-the-trainer course.[16]

Some organisations expected to be able to implement FRAIA processes themselves after the pilot. On the other hand, some organisations have not yet mastered all the material from the FRAIA and find external guidance and expertise important. Topics where guidance and expertise were desirable and frequently mentioned included:

- Interpretation of the FRAIA questions;
- Completeness of responses;
- Fundamental rights and public values;
- Severity assessment of affected fundamental rights;
- Time management.

The following are some examples of comments by the participating parties:

"It isn't always clear what is meant, exactly [by the questions]. The process supervisors were there to explain this time, but it is important to also have someone there in future who can lead the discussion and keep it on track."

"This is relatively new to us. You are not adequately versed in it to go deeper into things at the right time or to say that something can be skipped, so the calm and patient process supervision was very nice. Doing this on our own within the organisation is difficult: we would really need to train those people."

[16] See https://professionals.uu.nl/nl/cursus/impact-assessment-mensenrechten-en-algoritmes-iama.

"We would definitely do a FRAIA again. It is not always easy to judge whether something is a good case; we are bound to miss the mark on that at some point, but we will learn from that."

"Running a FRAIA independently may still be complicated in terms of how it is taken on within the organisation in practical terms. What are the ethical values? I'm not very familiar with that myself. It's a complicated subject matter. In that sense, an external supervisor would be good."

Not all participating organisations will directly integrate the FRAIA into their own algorithm-related work processes. The following distinction can be made:

- Some organisations are already going through the FRAIA where possible and necessary, depending on available capacity and risks related to the algorithm.

- Some organisations state that they would like to use the FRAIA as an accountability document (e.g., through the Algorithm Register) but that the FRAIA is too time-consuming in its current form. These organisations would like to see an abridged FRAIA document, possibly with a FRAIA quick scan and DPIA integration.

- Some organisations miss FRAIA's integration with laws and standards that dictate the requirements for algorithms. It should be emphasised here that the FRAIA is not intended to prescribe laws and standards but to promote interdisciplinary discussion across a wide range of domains. The Algorithm Framework may come closer to a tool or environment that prescribes laws and standards regarding the use of algorithms.[17]

In terms of future outlook, the above information shows that there will be organisations that are unable to conduct the FRAIA independently and want external guidance, that not all organisations have the capacity to implement FRAIA processes, and that the FRAIA is not attractive to all organisations in its current form.

On the other hand, it has also been compared to the GDPR/DPIA:

[17] See https://algoritmes.pleio.nl/groups/view/bf169271-70df-47b3-ae59-b46f6b1b32dc/algoritmekader.

"No one was thinking about privacy before, and that has transformed organisations as well. The same will happen again now [with the FRAIA]."

3. Conclusion

The publication of the FRAIA in 2021 created an opportunity to carefully map the intersection between fundamental rights and algorithms. Organisations are not required to do so, but the need is widely felt within the Dutch government. The FRAIA pilots were launched to encourage organisations to use the FRAIA for potentially high-risk algorithms. The goals of these pilots were (1) to provide support in building knowledge and expertise around FRAIA among participating government organisations and (2) to retrieve valuable feedback. This feedback led to numerous insights documented in the Findings section, which in turn formed the basis for the Recommendations section.

Participating organisations recognise the usefulness of a FRAIA process. Participants appreciate the interdisciplinary discussions, stating that the different insights contribute to improved visibility of the algorithms involved; most organisations consider the FRAIA PDF to be valuable as a reference book, and the process increases overall awareness of fundamental rights, ethics and algorithms.

However, the findings also show that some government organisations have yet to reach the necessary maturity level regarding algorithms to successfully complete a FRAIA process. They may not have enough staff or the right expertise, or there may be little to no support for or awareness of this type of process. At the same time, a number of organisations want to implement the FRAIA but argue that the FRAIA tool in its current form is still inadequate for their work practices in certain areas.

The FRAIA is receiving increasing political and government interest due to the political focus on algorithms and the recently approved AI Act. However, the pilots have shown that further explanation and adjustment of the FRAIA tool and its methodology is necessary before it can give substance to the political desire for a mandatory human rights assessment or the FRIA mentioned in the AI Act. The Recommendations section discusses this further.

4. Discussion

The FRAIA pilots have provided a lot of relevant feedback, with a few caveats:

- Participating organisations signed up of their own accord; participation was entirely voluntary and free of any obligation. This means that participating organisations already had some intrinsic motivation to work with the FRAIA. This may skew the feedback, which was generally very positive. By definition, these pilots are missing the view of an organisation that would not be open to the FRAIA or conducting a FRAIA process.[18]

- Because organisations signed themselves up to participate, they all a) had an appropriate case (potentially after consultation) and b) felt that their case was 'good enough', at least. It also reduced the odds of a very high-risk or controversial algorithm being tested.

- Fifteen government organisations gained experience with the FRAIA during these pilots. Other government organisations have also gained experience with the FRAIA in recent years. The approach and guidance provided by UU and RIG have been fairly similar throughout these pilot processes, while other processes may have taken a different approach. Those experiences are missing from this report.

- It should be emphasised that there is a lot of activity regarding the responsible use of algorithms and AI right now. In addition to the FRAIA, many review frameworks, assessments and tools have appeared in recent times that contribute to the responsible development of algorithms and AI, such as the Non-Discrimination By Design Handbook (Handreiking non-discriminatie by design)[19], the Algorithm Assessment Framework (Toetsingskader Algoritmes) by the Netherlands Court of Audit[20], the Algorithm Research Framework (Onderzoekskader Algoritmes) by the Central Government Audit Service[21] and Algorithm Framework for Responsible Deployment Of Algorithms

[18] In some cases, only the initiator in a participating organisation was enthusiastic, and the other participants did not share their enthusiasm. While feedback from those processes might say something about how other sceptical organisations would view the FRAIA, it is also not a one-to-one comparison. An enthusiastic initiator who provides a suitable case can already make a big difference compared to an organisation where no one is interested.

[19] See https://open.overheid.nl/documenten/ronl-3f9fa69c-acf4-444d-96e1-5c48df00eb3c/pdf.

[20] See https://www.rekenkamer.nl/onderwerpen/algoritmes-digitaal-toetsingskader.

[21] See https://www.rijksoverheid.nl/onderwerpen/rijksoverheid/documenten/rapporten/2023/07/11/onderzoekskader-algoritmes-adr-2023.

(Algoritmekader Verantwoorde Inzet Van Algoritmen) by the Interior Ministry[22]. In addition, the European AI Act was adopted in March 2024. Although the details of this legislation are not fully known at the time of writing this report, it will undoubtedly have an immense impact on this field of work and, thus, the outlook of the FRAIA. The instruments mentioned above and the upcoming European legislation all have a slightly different focus and objective, but there is a lot of overlap. Many government organisations are still struggling to determine which tool is most appropriate to apply to their algorithms and how the FRAIA and AI Act relate to each other and the other tools. To determine the role of the FRAIA for government organisations in future, these other developments cannot be disregarded, and the coherence of these tools must be analysed.

- The FRAIA is intended as an assessment or tool to determine whether fundamental rights are affected and to have a discussion about that. In practice, the FRAIA is also seen as a test to determine compliance with laws and regulations. This is a different approach than an assessment. It is important to keep the goals in focus, also in relation to the previous point in terms of consistency with other frameworks and new legislation.

[22] See https://www.rijksoverheid.nl/documenten/rapporten/2023/06/30/implementatiekader-verantwoorde-inzet-van-algoritmen.

5. Recommendations

This section contains numerous recommendations related to the FRAIA tool. Some of these recommendations substantiate the need for further development of the FRAIA.

5.1. Before initiating the FRAIA

1. For each case, estimate whether the FRAIA should be conducted. The authors recommend only using the FRAIA on high-risk algorithms.[23]

    a. Some organisations already have an algorithm checklist to determine whether an algorithm should be published in the Algorithm Register. Such a checklist can provide guidance in assessing whether to use the FRAIA.

    b. Developing a special FRAIA quick scan can help with this, albeit integrated with existing quick scans such as a pre-DPIA or information security quick scan.

2. For each case, determine at what point in the algorithm life cycle to deploy the FRAIA. In theory, all questions can be answered after development, but in practice, that will lead to the discussion of questions that have already been answered by a development team weeks, if not months earlier. We recommend developing guidelines for this purpose.

3. Ensure clear purpose identification. Establish ahead of time what the FRAIA will be used for, what is to be delivered and to whom, who owns or is responsible for the document, and whether the document needs to be kept current. Almost all participating organisations trialled an algorithm that had already been developed or deployed. The FRAIA was used to retrospectively test whether potential risks had been adequately mitigated or covered. This raises the question of whether it is better to complete the FRAIA before, during or after development. This should be stated more clearly in the FRAIA document.

4. Determine how the FRAIA compares to similar instruments. This prevents extra work and creates clarity about an organisation's needs at the relevant stage of the algorithm life cycle.

[23] The AI Act and the guideline on the Algorithm Register website provide a definition of "high-risk algorithms": https://www.digitaleoverheid.nl/overzicht-van-alle-onderwerpen/algoritmes/algoritmeregister/.

a. The pilots show that organisations struggle to determine when to use which tool in the development of algorithms. We therefore recommend it is described in the FRAIA how it relates to other Dutch and European instruments.

b. Conducting a FRAIA can help comply with legislation such as the AI Act but may not be sufficient in many cases. We recommend examining which legislation the FRAIA can fulfil and include this in the FRAIA for completeness.

5.2. Session structure

1. The sessions should be conducted on-site with as few online participants as possible. If necessary, an exception can be made here for external vendors or anyone listening in.

2. Avoid inviting too many people to a FRAIA session, but don't invite too few, either. On average, four to eight participants per session should be sufficient. Remember flexibility: you can always change the number of participants or roles before the start of the next session.

5.3. Process supervision

1. Involve a process supervisor to guide the discussion and progress.

2. Insofar as possible, this process supervisor should be substantively independent of the case to ensure objectivity.

3. A good process supervisor should emphasise the importance of customisation: no case will be the same, and each case deserves a unique approach with respect to session length, number of participants, and any substantive support.

4. We recommend investigating whether a government-wide trainer pool for FRAIA process supervisors can be set up by the Interior Ministry or the Association of Netherlands Municipalities, for example. This can be of particular value to smaller government organisations. As they deploy relatively few high-risk algorithms, the number of FRAIA processes they can expect to complete is limited, and training an in-house FRAIA supervisor may be too great an investment. A government-wide trainer pool could solve that.

5.4. FRAIA form and content

1. There are several aspects related to the form and content of the FRAIA that require review. A number of components received negative feedback, such as the job list of possible participants, the process in the interactive PDF, and the

lack of things like an action list. Other similar feedback is too detailed and falls outside the scope of this report. However, we would like to modify all sections that received negative feedback to minimise friction.

2. The FRAIA frequently refers to other tools, such as the DPIA or the Non-Discrimination by Design Handbook. We recommend integrating the FRAIA with the DPIA in a future version *in any case*. The vast majority of organisations stated that conducting a FRAIA and a DPIA is redundant, as questions frequently overlap.

3. The pilots show that some questions from the FRAIA can be interpreted in different ways. The relevant questions have been noted and we recommend making them more specific in a future version to ensure that the question means the same thing to every organisation.

4. In the authors' view, all parts of the FRAIA are necessary to arrive at an acceptable trade-off between the benefits of the algorithm and its impact on fundamental rights. We strongly discourage skipping entire sections.

5.5. After completing the FRAIA

1. Make sure the completed FRAIA document is stored in a central location where stakeholders can easily find it for reference.

2. Treat the FRAIA as a dynamic document: answers in the FRAIA can be modified or supplemented as the project progresses. This ensures that the algorithm documentation remains up to date.

3. Maintain version control and keep clear who owns and maintains the document. This consideration could also be more strongly embedded in the FRAIA itself.

5.6. Knowledge transfer

1. Consider how to better streamline FRAIAs between municipalities. Some municipalities may use the same algorithms. The question that arises is whether each municipality has to complete a FRAIA separately or whether things can be taken over in full or in part. The external vendor may be able to play a role in this.

2. The desire for a “FRAIA community” was frequently voiced during these pilots. Explore how best to accomplish this so that organisations can learn from each other and not have to reinvent the wheel each time.

3. Roll out the FRAIA even more broadly than the current pilot projects. This will help more widely disseminate knowledge and experience about and increase support for the FRAIA. One of the participating municipalities had the following to say to the Interior Ministry: *"Don't limit yourself to the eight processes currently in place, but roll it out more broadly! This will increase support among municipalities."*

5.7. Government, laws and regulations

1. To avoid confusion, it is important that the terms "fundamental rights" and "human rights" are used appropriately by politicians and governments. We recommend explaining these terms in the FRAIA and aligning them with their definitions in the AI Act.

2. The AI Act includes a requirement for a Fundamental Rights Impact Assessment (FRIA) for high-risk AI. Various parties consider the FRAIA to be one of the most advanced assessments that could potentially fulfil the FRIA requirement. As such, the FRAIA is receiving attention from other EU member states. We recommend exploring whether and how the FRAIA can be aligned with FRIA, albeit in cooperation with the recently established EU AI Office.

3. The AI Act also mandates a conformity assessment for high-risk AI systems. In practice, there appear to be misunderstandings about impact assessments, risk assessments, compliance assessments and conformity assessments. It is important to establish what the FRAIA is and what it is not. We recommend clarifying this explicitly in the Algorithm Framework.

4. Explore how the FRAIA can be integrated with the Algorithm Framework. This is also partly related to the conclusions related to overlap and cohesion (8, 18 and 19). Examples from other countries can serve as inspiration here, such as the HUDERAF framework.[24]

[24] https://www.turing.ac.uk/news/publications/human-rights-democracy-and-rule-law-assurance-framework-ai-systems.

6. Acknowledgements

Conducting the pilots and achieving the goals would not have been possible without the enthusiasm, dedication and commitment of the participating parties. The authors would like to thank the participating organisations for their openness and willingness to contribute case studies and share feedback: the Tax and Customs Authority, the Municipality of Almere, the Municipality of Assen, the Municipality of Den Haag, the Municipality of Enschede, the Municipality of Haarlem, the Municipality of Den Bosch, the Municipality of Veenendaal, the Royal Netherlands Marechaussee, Logius, the Education Inspectorate, the Open State Foundation, the Province of Noord-Brabant the Province of Zuid-Holland.

Special thanks go out to the process initiators who sparked enthusiasm within their organisations and were always pleasant partners with whom to set up and conduct the FRAIA sessions.

Lastly, the authors would like to thank the participating third-party vendors for their willingness to share their knowledge and insights about the algorithms. Their participation allowed the successful completion of the FRAIA and enabled all parties to be served.

7. Appendix

7.1. Method

The Method section is included for interested organisations that want to get started with the FRAIA and need a guideline for setting up FRAIA sessions. This section is divided into two parts because RIG and UU took different approaches to the external supervision of the FRAIA pathways in some cases. This was related to available capacity, the varying composition of supervisors and the two organisations' varying roles and backgrounds. Both methods are described below. The processes were conducted throughout the year 2023.

7.1.1. Utrecht University

As the author of the FRAIA, UU has been gaining experience with the tool since 2021. This first happened during the test phase in early 2021, during which the FRAIA was tested with four government parties (the Municipality of Amsterdam, the Municipality of Rotterdam, the National Office for Identity Data and the Tax and Customs Administration). After the publication of the FRAIA, UU supervised several other FRAIA processes, including for the Municipality of Rotterdam. A standard format emerged from these experiences, which UU uses when supervising FRAIA processes. The format is as follows:

Type of meeting	Duration (h)	FRAIA components	On-site or online
Intake	0.5	-	Online
First FRAIA session	2	Part 1, Part 2A	Preferably on location
Second FRAIA session	2	Part 2B, Part 3	Preferably on location
Third FRAIA session	1	Part 4	Preferably on location
Evaluation	0.5	-	Online
Total:	6		

Table 3 Representation of a standard FRAIA process by the UU.

The FRAIA generally has a standard format, with some deviations by UU in certain cases due to availability, scheduling or case complexity. FRAIA sessions can also be conducted entirely online. On-site is preferred, as it promotes discussion and speeds up the process (see also Findings).

Two supervisors supervised all eight processes. UU had a pool of three experienced FRAIA supervisors available with different areas of expertise (digital ethics, philosophy, human rights, and *critical data studies*). During the FRAIA sessions, one

supervisor served as a moderator, introducing questions, making sure answers were written down correctly in the FRAIA document, and providing explanations in the event of ambiguity about certain concepts or questions. The other supervisor took notes (in scientific terms: *field notes*) of the process: what questions were perceived as difficult, where were the biggest ethical sticking points, who (which functions) were present during the sessions, and what role they played in the discussions.

The FRAIA sessions were mostly held on-site at the government organisation in question. The government organisation invited its own people to the sessions based on the table of required functions in the FRAIA (p. 2 in the FRAIA document). One of the attendees was asked beforehand to take minutes during the sessions. The document was then reviewed jointly, with the minute taker typing directly along in the FRAIA.

Answers are immediately visible to all participants, making it easy to add or correct answers until everyone agrees and can move on to the next question. At the end of the process, the completed FRAIA document is stored on a shared drive and updated as necessary in the future as the algorithm or use thereof changes. In principle, UU does not have access to the completed FRAIA document. The relevance of the process lies in the feedback gathered from it; in terms of the content of the case, the process is for the organisation in question.

Online evaluations were conducted based on a standardised questionnaire after the FRAIA sessions. This questionnaire was compiled in advance with feedback from all parties (UU, RIG, Interior Ministry) to ensure that the evaluators collected the same feedback and that the feedback was useful to the principal. RIG also used the questionnaire, which can be found in the Appendix.

7.1.2. Rijks ICT Gilde

The Rijks ICT Gilde (RIG) is part of the Ministry of the Interior and Kingdom Relations and falls under the National Organisation for Development, Digitalisation and Innovation. RIG staff are seconded throughout the national government to provide support on technology issues. RIG has a strong focus on the responsible use of data, algorithms and AI, working with various government organisations on this topic. Several RIG employees have already gained experience with FRAIA processes, including with Toeslagen (benefits) and the Judiciary. In preparation for the start of the pilots, relevant RIG staff also attended the UU “train-the-trainer” workshop. A RIG employee also participated in the FRAIA sounding board group on behalf of the RIG. RIG uses the following format for the FRAIA pilots:

Type of meeting	Duration (h)	FRAIA components	On-site or online
Intake	0.5	-	Online
First FRAIA session	2	Part 1	Preferably on location
Second FRAIA session	2	Part 2A, Part 2B, Part 3	Preferably on location
Third FRAIA session	1	Part 4	Preferably on location
Evaluation	0.5	-	Online
Total:	6		

Table 4 Representation of a standard FRAIA process by RIG.

It should be mentioned that this is the standard format, which was deliberately deviated from in some cases due to the time available in the organisation, the availability of RIG staff or the complexity of the case. Different formats have been tried for different processes. For example, the first session also covered Part 2A of the FRAIA several times, and the third session covered Part 3 and Part 4 in two hours.

The intake always took place online using an intake form (see Appendix). During the intake, RIG explained the FRAIA and the format of the various sessions; the organisation then set up its own case study. This allowed RIG to assess the suitability of the case for the pilot and to focus the sessions on the particular case in the preparations.

Participating organisations were asked to have all employees involved attend in person, as online attendees are less able to engage in an in-person discussion. Some FRAIA sessions took place in a hybrid format due to circumstances, with some participants participating online.

Three RIG employees were usually involved in each pilot: one for substantive or content guidance, one for process guidance, and one to evaluate the FRAIA as a whole (in connection with writing the current report). RIG had a pool of six employees available. These employees have diverse backgrounds, including data science, mathematics, artificial intelligence, computer science and philosophy. During the sessions, the content officer presented the content of the session and asked critical questions when the organisation answered the questions from the FRAIA. The process officer kept track of time, ensured that everyone had a say, and made sure the answers to the questions were written down correctly.

The sessions started with a brief presentation on the content of the particular session. During the first FRAIA session, RIG gave a brief introduction to the FRAIA. The participating organisation also gave a brief introduction about the algorithm selected

as a case study. Participants then began filling in the FRAIA. The FRAIA document was displayed on a large screen with the answers written out in the document by the minute taker. The pilots alternated between a minute taker from RIG and one from the organisation itself to determine the optimal arrangement. After the sessions, the organisation was responsible for further development of the FRAIA document. If RIG had taken the minutes, it offered recommendations for further development.

Finally, an evaluation took place using a standardised questionnaire. This questionnaire was compiled in advance with feedback from all parties (UU, RIG, Interior Ministry) to ensure that the evaluators collected the same feedback and that the feedback was useful to the principal. In some cases, the evaluation took place immediately after the last session; other times, it was conducted online afterwards. The questions in the questionnaire were related to the FRAIA process, and the answers to these questions can help both RIG and the respective organisations prepare for future sessions. The output from the various evaluations also provides support for this report.

7.2. The role of the moderator

During the FRAIA pilots, UU and the RIG were there to guide the organisations through the document. The guidance provided by UU is primarily process-oriented in nature: substantively, UU does not interfere with the answers unless, by way of exception, substantive guidance is needed. However, some processes had some substantive preparations done by the supervisors before starting Part 4 of the FRAIA. This part is often perceived as complicated. This additional substantive preparation allowed the supervisors to focus thoroughly on making sure the group arrived at the correct fundamental rights in Part 4. The same is true for RIG, with the difference that RIG was willing to provide substantive guidance for all components if requested. This could include technical concepts or information about the FRAIA in the context of the government landscape.

Feedback during the evaluations indicated that participating organisations appreciated the external guidance, which was perceived as pleasant and important. Directing, summarising, and sometimes ending discussions were seen as effective in keeping the session on track. The moderators felt comfortable operating from an outside view, as they felt free to park a discussion for a moment or to ask basic or critical questions on a topic that was taken for granted by the participants themselves. This endorses the substantive independence of the moderators. By having no substantive interest in the case, counselling can be as objective as possible. There is also no tendency to conform to the group or certain individuals.

7.3. Overview of all roles present at the FRAIA processes

This table can serve as inspiration for other government organisations regarding the types of roles that can be invited to a FRAIA process. The list is not exhaustive, and duplicate roles have been omitted. For example, the project manager was almost always present in each process, but the position of "project manager" appears only once in this table.

Roles attending the FRAIA processes
Business Operations Consultant
CIO Office Advisor
Data Ethics Consultant
Department Head
General Lawyer on Legal Affairs team
Analyst
Policy Advisor
Policy Officer
Citizen
Business Analyst
CDO
CIO
CISO
Communications Advisor
Consultant [external vendor]
Contract Manager
Information Security Coordinator

Data Analyst
Data Engineer
Data Engineer [external vendor]
Data Manager
Data Scientist
Deep Learning Developer [external vendor]
Executive Secretary
DPO and Compliance Officer [external vendor]
Data Protection Officer
User of the algorithm
Head of Research
Information Analyst
Information Manager
Legal Advisor
Legal Advisor, GDPR
Legal Advisor, privacy
Legal Expert
Environmental Director
Algorithm Developer
Client
Founder [external vendor]

Participation Coach
Privacy Advisor
Privacy and Security Officer
Privacy Officer
Product Developer [external vendor]
Programme Manager
Project Leader
Project Manager, ICT
Project Secretary
Senior Communications Consultant
Senior employee [external supplier]
Strategic Information Manager
Science Officer

Table 5 All roles present during the FRAIA processes.

7.4. Standardised questionnaire for evaluation sessions

UU and RIG used this questionnaire to collect feedback from the participating parties in the FRAIA pilots.

Overall impression:

- *What did you think of the FRAIA process?*
- *What are the most important and useful insights you have gained?*
- *What is the nicest aspect of the FRAIA?*
- *What went well?*

- *What could have gone better?*

Clear areas for improvement:

- *Do you feel there are any pitfalls in your project that the FRAIA was not able to overcome or detect?*
- *Are there any important questions or considerations that are not sufficiently addressed in the FRAIA?*
- *What would you improve on the FRAIA (the tool itself or the process)?*

Future vision:

- *Would you complete a FRAIA again in the future for the same project in the event of significant changes to the project, for example?*
- *Would you use the FRAIA for another project in the future?*
- *Do you feel you could use the FRAIA independently now? If not, what do you need to be able to do so?*

Implementation:

Substantive:

- *Were the FRAIA sessions properly prepared? Please explain your answer.*
- *What will you do with the results? How and with whom will the results be communicated? Do you need any help with that?*

Procedural:

- *At what stage(s) of a project do you think the FRAIA would be most useful for your organisation?*

- *Does your organisation have any way to do some sort of "quick scan" in advance to avoid wasting time on an unnecessary FRAIA process for algorithms that ultimately turn out not to be high-impact?*

- *What type of role or background do you think should absolutely be present at the various FRAIA sessions? Who is indispensable at the first, second and third meetings (e.g., project leader, privacy officer, data scientist, etc.).*

7.5. RIG intake form

The intake checklist below was used to decide whether an organisation could participate in the pilot. RIG also used this checklist to prepare for the sessions (e.g., predict potential bottlenecks, etc.).[25]

1. ***Reason***

    a. *Why does your organisation want to participate in the FRAIA pilot?*

    b. *What do you hope to get out of participation in the FRAIA pilot, and what do you consider a successful outcome of the FRAIA pilot for you?*

    c. *What is your greatest concern?*

    d. *What do you expect afterwards?*

2. ***Algorithm***

    a. *Which algorithm is concerned?*

    b. *What is the algorithm for?*

    c. *What stage is the algorithm in (planned/in development/production)?*

    d. *Can you provide more information (in outline) on:*

        i. *Algorithm input*

        ii. *Type of algorithm*

[25] Although UU also conducted intake interviews with organisations to investigate the party's motivation and the suitability of the case presented, no standardised form was used for this purpose.

iii. *Algorithm output*

e. *What documentation is available for the algorithm (and can it be shared with RIG)?*

3. FRAIA pilot structure

a. *What the FRAIA pilot offers:*

i. *An introduction to a methodology (FRAIA) to support responsible deployment of AI (like Prince-2 does for project management)*

ii. *An opportunity to gain experience with this methodology in order to apply it in the organisation itself in future*

iii. *Starting points for further research regarding the algorithm*

b. *What the FRAIA pilot does not offer:*

i. *Participation in the FRAIA pilot does not lead to any certification*

ii. *RIG does not issue a final ruling or advice on the algorithm (but can share observations and tips with the organisation if desired)*

iii. *A full, in-depth FRAIA will not be completed during the pilot*

4. FRAIA pilot process

a. *Intake*

b. *Record of agreements + proposal (RIG)*

i. *Point of contact/project owner at the organisation*

ii. *3 sessions: 1–2 weeks between sessions*

iii. *4 parts; specifics in consultation*

iv. *Participants per section (see next page)*

c. *Three FRAIA sessions on site*

i. *Preparation*

1. *Logistical preparation (organisation)*
2. *Process preparation (RIG)*
3. *Content preparation (participants)*

ii. *Guidance/support*

1. *1 process supervisor (RIG)*
2. *1–2 substantive support (RIG)*
3. *1 minute taker (organisation)*

d. *Evaluation*

i. *Issue questionnaires to participants (organisation)*

ii. *Summarise responses (organisation)*

iii. *Discuss evaluation (on-site or online meeting)*

5. Intended results of FRAIA pilots

a. *Completed FRAIA (RIG + organisation)*

b. *Evaluation of FRAIA (organisation, based on RIG questionnaire)*

c. *Final report based on evaluations (for Interior Ministry)*

d. *Academic research based on evaluations (UU)*