Algorithm Audit

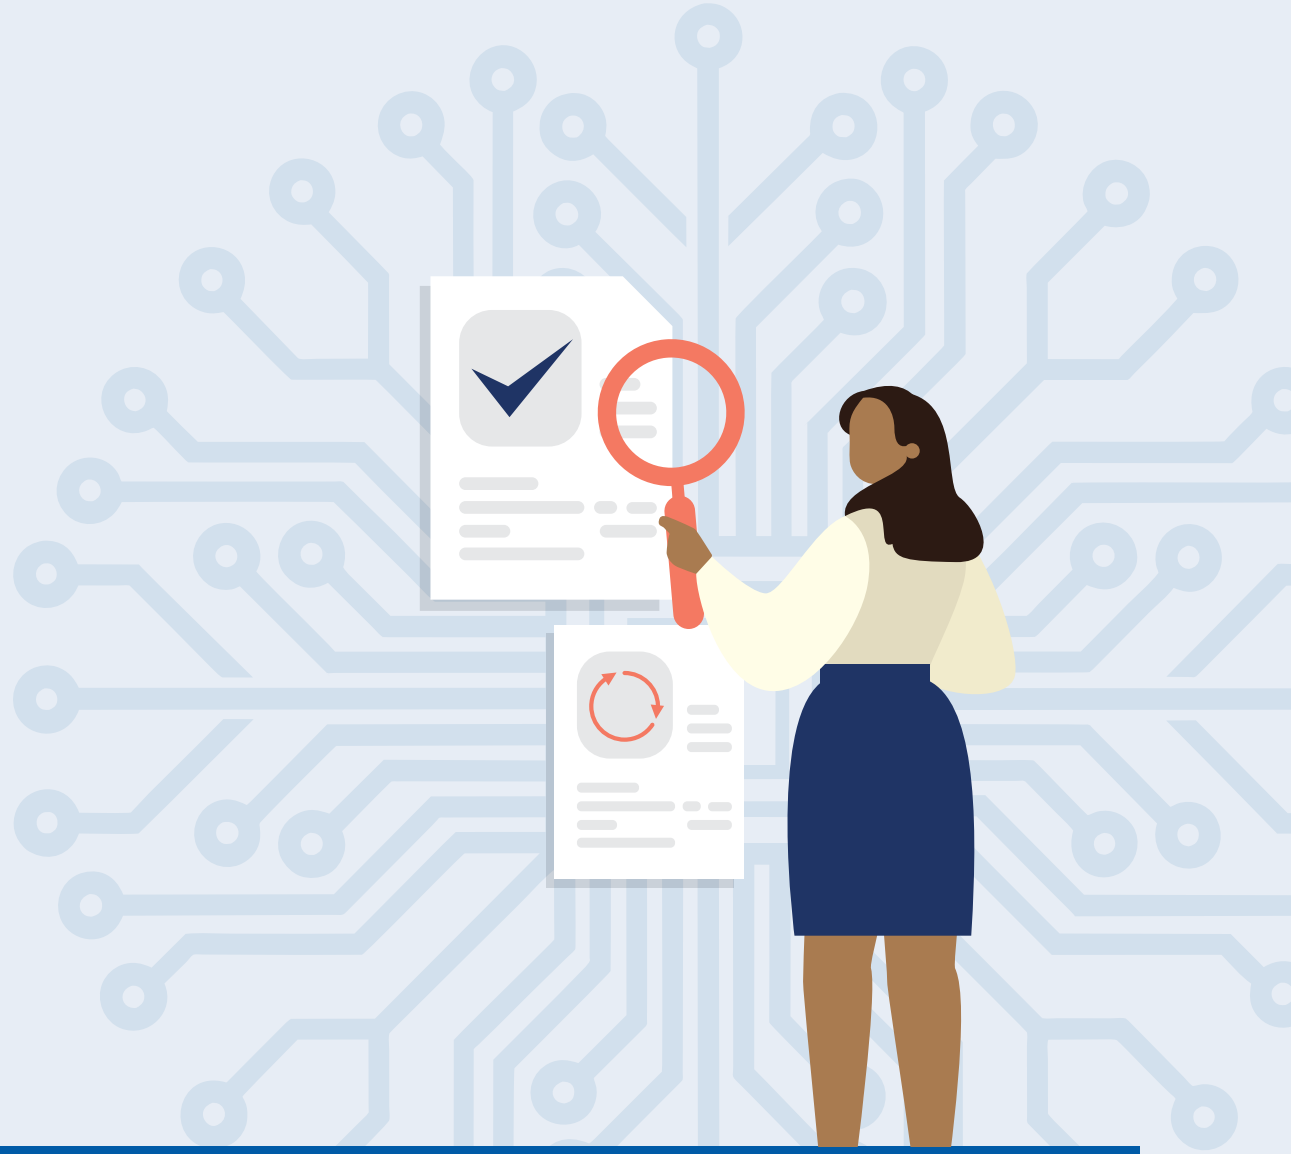

AI AQT documentation

Open source AI and Algorithm Qualification Toolkit for EU and Dutch policy frameworks

May 2026

Table of Contents

1. **Introduction** 3
2. **Questionnaire Identification – AI system** 6
   - 2.1 Inference 7
   - 2.2 Autonomy 7
3. **Questionnaire Identification – Solely automated decision-making** 13
   - 3.1 Decision, legal or similar effects and human involvement 13
4. **Questionnaire Identification – High-impact algorithms** 24
   - 4.1 Direct consequences 26
   - 4.2 Significant effect on the outcome of the process 26
5. **Questionnaire Role and status** 35
   - 5.1 Operator roles 35
   - 5.2 Status of the AI system 36
6. **Questionnaire Risk category** 39
   - 6.1 Annex I – List of Union harmonisation legislation 39
   - 6.2 Biometrics branch 41
   - 6.3 Transparency obligations and prohibited non-consensual fake nude imagery 43
   - 6.4 Annex III high-risk domains 44
   - 6.5 Article 5 horizontal prohibitions 49
   - 6.6 Article 6(3) – Profiling and limited tasks 50
   - 6.7 Article 2 exception 51
7. **Questionnaire Obligations** 53
   - 7.1 Obligations for prohibited AI systems 53
   - 7.2 Obligations for high-risk AI systems 53
   - 7.3 Obligations for generative and interactive AI 55

1. Introduction

The AI and Algorithm Qualification Toolkit (AI AQT) helps with navigating policy frameworks relevant when applying algorithmic systems, such as the GDPR, the AI Act and national policies. The European Union's (EU) General Data Protection Regulation (GDPR) is a cornerstone framework governing how algorithms process data. It regulates how organizations collect, use and share individuals' personal data – whether through analogue means or algorithmic. The EU AI Act has introduced requirements for artificial intelligence (AI) systems to safeguard the safety, health and fundamental rights of EU citizens. In addition, some countries, such as the Netherlands, have implemented broader control measures for algorithmic systems that may have severe impacts on stakeholders. Navigating which legal frameworks apply to a given algorithmic system, and remaining compliant with limited resources, can be challenging. The AI AQT serves as a building block toward compliance across these various policy instruments.

The AI AQT consists of multiple user-friendly, dynamic questionnaires designed to support the harmonized identification and risk classification of algorithmic systems.

Questionnaire **Identification** determines which laws and regulations your algorithmic application falls under:

- **AI system** – As defined in the Article 3 AI Act.
- **Personal data** – Whether the General Data Protection Regulation (GDPR) is applicable.
- **Profiling** – As defined in Article 4(4) GDPR.
- **Solely automated decision-making (sADM)** – Automated decision-making (ADM) practices, including profiling, as restricted or prohibited under Article 22 of the GDPR.
- **Optional:** **High-impact algorithm** – Systems with severe impact on stakeholders monitored by the Dutch government (irrespective of their status under the AI Act). See Algorithm Register Guidelines.

General Purpose AI (GPAI) and requirements for GPAI model providers, as set out in Articles 51-55, fall outside the scope of the AI AQT.

Questionnaire **Role and status** determines what role you have with respect to the AI system and what the usage status is:

- Role:
  - **Provider** – Natural or legal person that develops an AI system or general-purpose AI model and places it on the market under its own name or trademark (Article 3(3) AI Act).
  - **Deployer** – Natural or legal person using an AI system under its own authority in a professional context (Article 3(4) AI Act).
  - **Authorized representative** – A natural or legal person located or established in the Union who has received and accepted a written mandate from a provider of an AI system to carry out on its behalf the obligations and procedures established by the AI Act (Article 3(5) AI Act).
  - **Importer** – Natural or legal person established in the Union who places on the market an AI system bearing the name or trademark of a non-EU entity (Article 3(6) AI Act).
  - **Distributor** – Natural or legal person in the supply chain that makes an AI system available on the Union market without altering its properties (Article 3(7) AI Act).
- Status:
  - **In use** – Must comply by 2030 at the latest with requirements of the AI Act (Article 111 AI Act).

- **In development** – Shorter deadlines depending on the type of AI system (Article 113 AI Act).
  - 2 December 2027 (Annex III – high-risk domains)
  - 2 August 2028 (Annex I – high-risk due to safety component in regulated product)
  - 2 December 2026 (AI generated content labelling)

Questionnaire **Risk category** determines the risk classification of the AI system:

- **No requirements** – AI systems with no requirements according to the AI Act (Article 5-6 AI Act).
- **Prohibited AI systems** – AI system meets definition of a prohibited AI system (Article 5 AI Act).
- **High-risk AI systems** – AI system that requires additional control measures (Article 6 and Annex I and III AI Act).
- **Generative and interactive AI** – As described in the Article 50 AI Act (transparency requirements).
- **Exception** – An exception applies for some AI system applications, among others, scientific research, defense, national security etc. (Article 2 AI Act).

The flow of the AI AQT is illustrated in Figure 1. Because the qualification criteria for AI systems, sADM and high-impact algorithms overlap, users are asked a single set of questions in Questionnaire Identification. As high-impact algorithms are particular of interest for Dutch public sector organisations, users can select whether they want this type of algorithms to be identified, or not. Background logic orchestrates the question flow and determines when enough information has been gathered to assess whether a system qualifies as an AI system, sADM and/or high-impact algorithm. Figure 2 presents a Venn diagram illustrating the possible outcomes from Questionnaire Identification and Risk classification. Algorithm Audit has also published supplementary materials on interpreting the AI Act's definition of an AI system[1], as well on the scope of sADM and the requirement of meaningful human intervention.[2] All AI AQT materials on the AI Act are updated according to AI Omnibus amendments.

The AI AQT is designed for organisations seeking a standardized and harmonized approach to classifying algorithmic systems. The tool offers a user interface suitable for both legal expert and users without legal background, translating legal concepts into accessible language.

The classification of systems into the respective (legal) categories follows a "better safe than sorry" principle, where a conservative classification is preferred over prematurely concluding a system is out of scope. The tool is intended for guiding further compliance actions and does not substitute professional legal assessment or consultation.

[1] Implementation of the AI Act – Definition of an AI system, Algorithm Audit (2025).
[2] Meaningful human intervention for risk profiling algorithms – Preventing decision-making based solely on profiling, Algorithm Audit (2025).

Open source code repository

The first version AI and Algorithm Qualification Toolkit (AI AQT) was developed in collaboration with the Municipality of Amsterdam. The tool's source code is available on GitHub and can be (re)used under the EUPL-1.2 license.

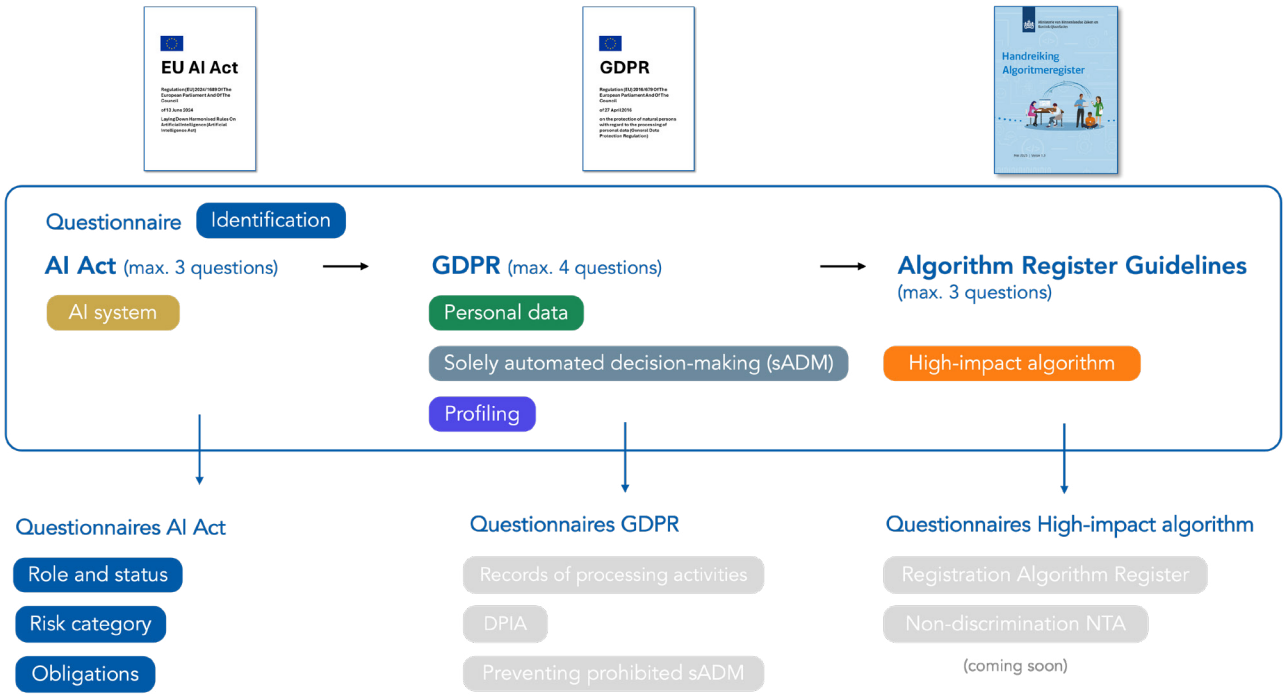

Figure 1 - Overview of the flow of dynamic questionnaires in the AI AQT. Responses to earlier questions are taken into to account in order to minimize the number of questions presented to the user.

This white paper details the considerations and design choices made during development of the AI AQT. It first explains the components of Questionnaire Identification used to determine whether an algorithmic system qualifies as an AI system under the AI Act (section 2). It then addresses GDPR aspects: determining whether personal data is processed, and whether profiling and sADM are at stake (section 3). Next, the concept of high-impact algorithms is introduced and it is discussed how such systems are identified (section 4). Thereafter, the role in relation to AI systems under the AI Act, depending on their usage status, is elaborated on (section 5). Thereafter, it is described how Questionnaire Risk category qualifies AI systems as prohibited or high-risk and whether an exception applies (section 6). The section concludes with obligations that apply for certain combinations of role, status and risk category (section 7).

Box 1

Reserved compliance with the AI Act and Article 22 of the GDPR

This document reflects Algorithm Audit's interpretation of the legal texts of the AI Act and Article 22 of the GDPR, relevant guidelines issued by the European Commission (EC) and European Data Protection Board (EDPB), as well as existing case law. No rights can be derived from this analysis. Ultimately, the competent courts determine the correct interpretation of what constitutes a (high-risk) AI system and sADM.

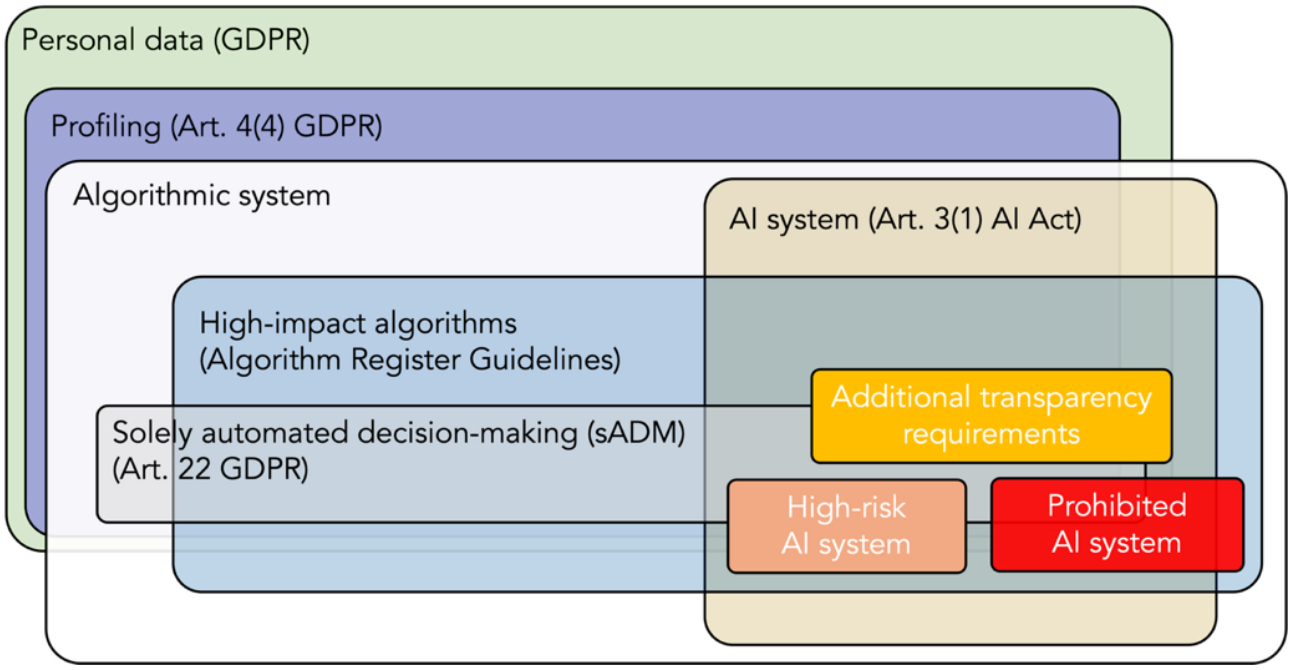

Figure 2 - Overview of the type of algorithmic systems identified by the AI AQT.

2. Questionnaire Identification – AI system

In Questionnaires Identification, Q1-Q3 focus on the definition of an AI system, i.e., the system's design and output. Q4-Q7 determine whether personal data is processed, whether profiling is at stake and whether data is stored and shared. Q8-Q12 determine whether sADM and/or a high-impact algorithm is used.

The definition of an AI system (hereafter: "AI definition") is set out in Article 3(1) of the AI Act. Only systems that meet this definition fall within the scope of the regulation. Article 3(1) defines an AI system as follows:

"… a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments."

As elaborated on in the Algorithm Audit's analysis "Implementation of the AI Act – Definition of an AI system"[3], only two concepts in this definition are central to distinguish AI systems from other algorithmic systems:

1. Inference
2. Autonomy.

The other parts of the AI definition are either optional characteristics, such as 'adaptiveness', or concepts that do not contribute to differentiating AI systems from other IT systems, such as 'machine-based'.

This section begins with an explanation of what is meant with inference (2.1) and autonomy (2.2). Guidelines provided by the EC (hereafter:

[3] Implementation of the AI Act – Definition of an AI system, Algorithm Audit (2025).

'guidelines') on the interpretation of the AI definition are incorporated in this analysis.[4] Next this section discusses how these concepts are incorporated into the AI AQT Questionnaire Identification as Questions 1-3 (Q1-Q3).

Algorithm Audit has published examples and explainers clarifying what type of outputs, inference and rule-bases systems fall within, and which systems fall outside the scope of the definition.[5]

2.1 Inference

The capability to infer is the most important element of the definition that distinguishes AI systems from other algorithmic systems. Recital 12 of the AI Act states that: *"A key characteristic of AI systems is their capability to infer"*. The guidelines state that inference is an *"indispensable condition"*. Inference means both the capability to derive models or algorithms from data and to derive outputs from input. Recital 12 further clarifies that (at least) two techniques enable inference: machine learning and logic- and knowledge-based approaches.[6]

There are many techniques which are not traditionally considered to be machine learning, such as statistical approaches where model parameters are 'fitted' on data. These techniques can still be considered inference from data, they are thus included in the AI definition. This is captured Q2.

Logic- and knowledge-based approaches to AI are different. Logic- and knowledge-based approaches are usually carefully manually designed. Although with these approaches there may not be a model derived from data, they are still considered as techniques that enable inference for their reasoning capacity. The question whether the application is a logic- and knowledge-based system is included in Q3.

The AI definition mentions "predictions, content, recommendation, or decisions" as forms of output of an AI system. Because deriving these types of outputs is conditional to conclude that the system has the capability of inference, the output type in included in Q1. Because the target audience of AI AQT – developers, product owners, line managers and other executive users – are typically familiar with the output of their application, Questionnaire Identification starts with asking about the output an algorithm produces.

More information on the concept of inference can be found in section 3 of Algorithm Audit's analysis "Implementation of the AI Act – Definition of an AI system".[7]

2.2 Autonomy

The one other concept in the definition required to distinguish AI systems from other algorithmic systems – aside from inference – is autonomy. Recital 12 of the AI Act indicates this to mean: *"AI systems are designed to be operate with varying levels of autonomy, meaning that they have some degree of independence of actions from human involvement and of capabilities to operate without human intervention"*. 'A certain degree' is however a weak requirement: a system does not have to be completely autonomous to meet this requirement. The guidelines mention that a system which generates an output by itself from manually provided in puts is already considered 'some degree of independence'.

This implies that every system that derives outputs itself is autonomous to a certain extent. In other words, if the inference requirement is met, the autonomy requirement is also met. We conclude that, in conjunction with inference, the 'autonomy' requirement imparts no meaningful criteria to

4 'Guidelines on the definition of an artificial intelligence system established by AI Act', European Commission (2025).
5 Examples and explainers AI Act, Algorithm Audit (2025).
6 Recital 12 of the AI Act.
7 Supra note 3.

distinguish AI systems from other algorithmic systems.

Identification Q1 – What type of output does the application derive?

The output generated by an algorithmic system gives an indication whether it qualifies as an AI system. Users of AI AQT are therefore first asked to indicate what type of output is generated by the system. See Figure 3.

As discussed in 2.1 Inference, deriving outputs is considered to be integral to the AI definition. Types of outputs should be predictions, content, recommendation, or decisions.

Prediction is a broad concept with differing interpretation across domains. In data science, a prediction does not have to be about the future. It can also relate to a data point that has not been observed before. In fact, every score, ranking, recommendation, label, classification, decision and generated content (image, text, speech etc.) is a prediction. Therefore, the answer options include explanatory terms for various types of predictions that should be recognizable to users (score, ranking, label, object-, face- or voice recognition). This choice favours accessible language. If one of these options are selected, users are forwarded to Question 2 (Q2). Note that multiple options can be selected.

Because dashboards are a common type of data-driven application, confusion about whether they qualify as an AI system can arise. On their own, dashboards only provide data visualization. No inference is at stake and they do not transcend *"the elementary processing of data by enabling learning, reasoning or modelling"*[8] as an AI system would. In and of itself, a dashboard cannot constitute an AI system, even if its coupled to one. If a user indicates the only output of a system is a dashboard, it is concluded it is not an AI system. Users are explicitly prompted to consider if other types of outputs are displayed in this dashboard. If they select one of the specified outputs along with "Dashboard", they are brought to Q2.

[8] Recital 12 of the AI Act and section 3 of Implementation of the AI Act – Definition of an AI system, Algorithm Audit (2025).

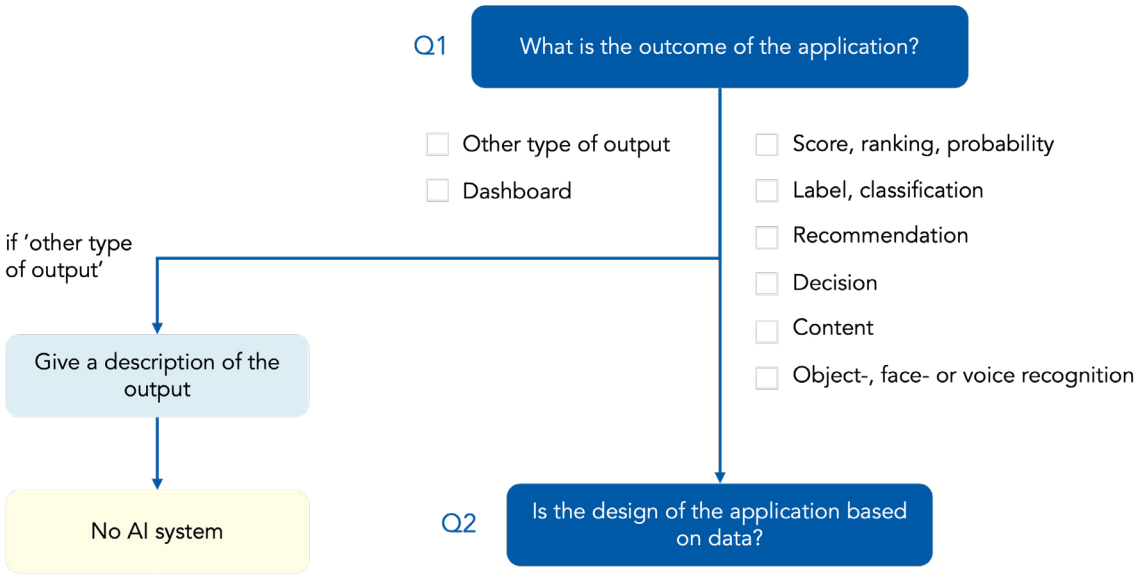

Figure 3 - Identification Q1 asks what type of output the algorithmic application generates.

The same logic applies to the option "Other type of output". If it is the sole provided answer, the tool concludes the application is not an AI system. If coupled with another output (other than dashboard), the user proceeds to Q2. In either case, the user is asked to provide a description of the output, which can be manually assessed by experts.

Identification Q2 – Is the design of the application based on data?

While Identification Q1 focuses on the system's output, Q2 examines how the output is generated (see Figure 4). Under the AI Act, an AI system is defined by its capacity to perform inference. As discussed in 2.1 Inference, determining whether the design of a system relies on data is a key factor in assessing whether inference takes place. This aspect is incorporated into Q2.

If the application contains components derived from data, then it is an AI system. This is the case, for example, when a model or algorithm is learned or fitted using statistics, optimization, simulation or machine learning or a similar technique. In this case, the conclusion – that the algorithm qualifies as an AI system – is presented. The user is then asked whether they want to continue with the rest of Questionnaire 1, which deals with 3. Questionnaire: Solely automated decision-making and 4. Questionnaire: High-impact algorithms, starting with Q8.

Recital 12 of the AI Act indicates that not all applications whose components are derived from data qualify as an AI system. Where design choices are made manually and data analysis is used only to inform those choices, the resulting algorithmic system does not constitute an AI system.[9] When this answer is selected, users are asked to provide an explanation. This information supports a case-by-case assessment, based on expert knowledge, of whether the algorithm does qualify as an AI system. After the clarification question, users are asked whether they want to continue with the rest of Questionnaire 1, which deals with 3. Questionnaire: Solely automated decision-making and 4. Questionnaire: High-impact algorithms, starting with Q8.

[9] This exception is discussed in detail in the paper "Implementation of the AI Act – Definition of an AI system". Supra note 3.

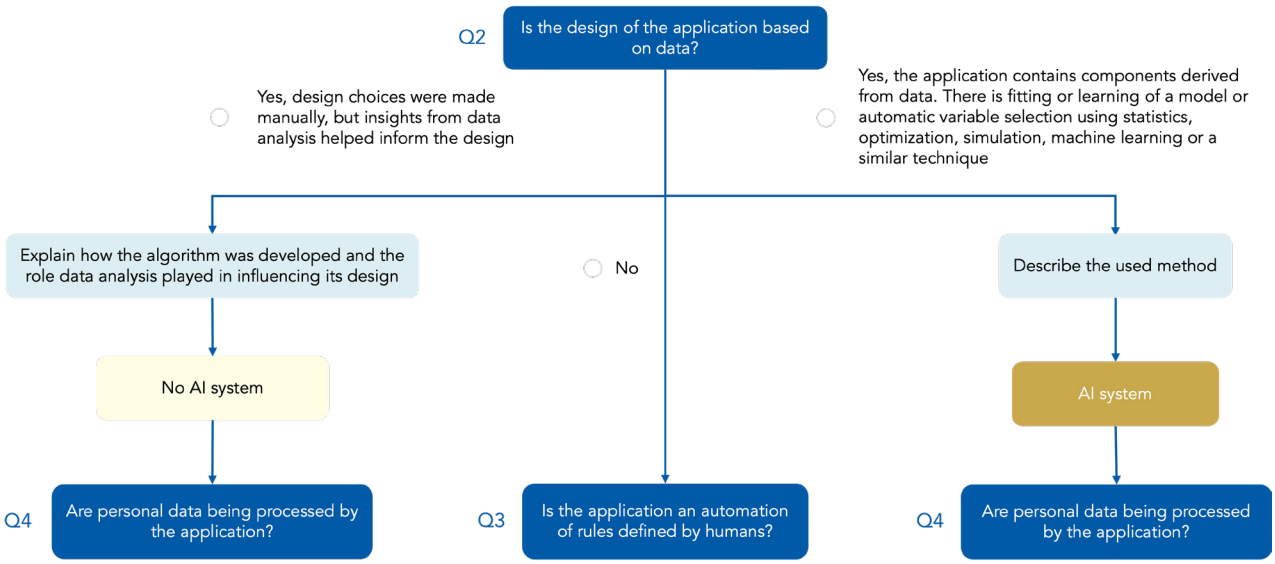

Figure 4 - Identification Q2 is about whether the design of the application is based on data.

Even if the design of the application is not based on data, the application can still be an AI system. To check if this is the case, users are redirected to Q3.

To assist users, the following remark is provided:

Data includes all forms of electronic information. Text, images, and audio are also data.

Applications can be designed manually. However, even when they are manually designed, the design is sometimes based on data analysis. For example, threshold values for (eligibility or exclusion) rules may be calculated from data, or criteria may be selected on the basis of calculated correlations.

In other cases, components (e.g., models and algorithms) are derived more automatically from data. This may involve, for example, fitting a statistical model to data or using machine learning to train a model or rule-based algorithm from data. Simulation and optimization techniques may also be used to derive a model from data.

Large language models such as ChatGPT are also derived (trained) from large amounts of textual data.

Identification Q3 – Is the application an automation of human-defined rules?

Q3 serves to capture a specific case in which no model is derived from data, but the system still qualifies as an AI system, following recital 12 of the AI Act and the guidelines provided by the EC. See Figure 5.

As discussed in 2.1, there exists a case where the design of a system is not based on data but there is still inference involved - that of case of logic and knowledge-based systems. These are often understood or referred to as rule-based systems that involve a certain level of manual programming. For this reason Q3 deals with how the rules a system follows are designed.

Logic- and knowledge-based approaches are explicitly stated in the AI Act as techniques that enable inference and thus should be considered AI. In practice these techniques are usually applied in

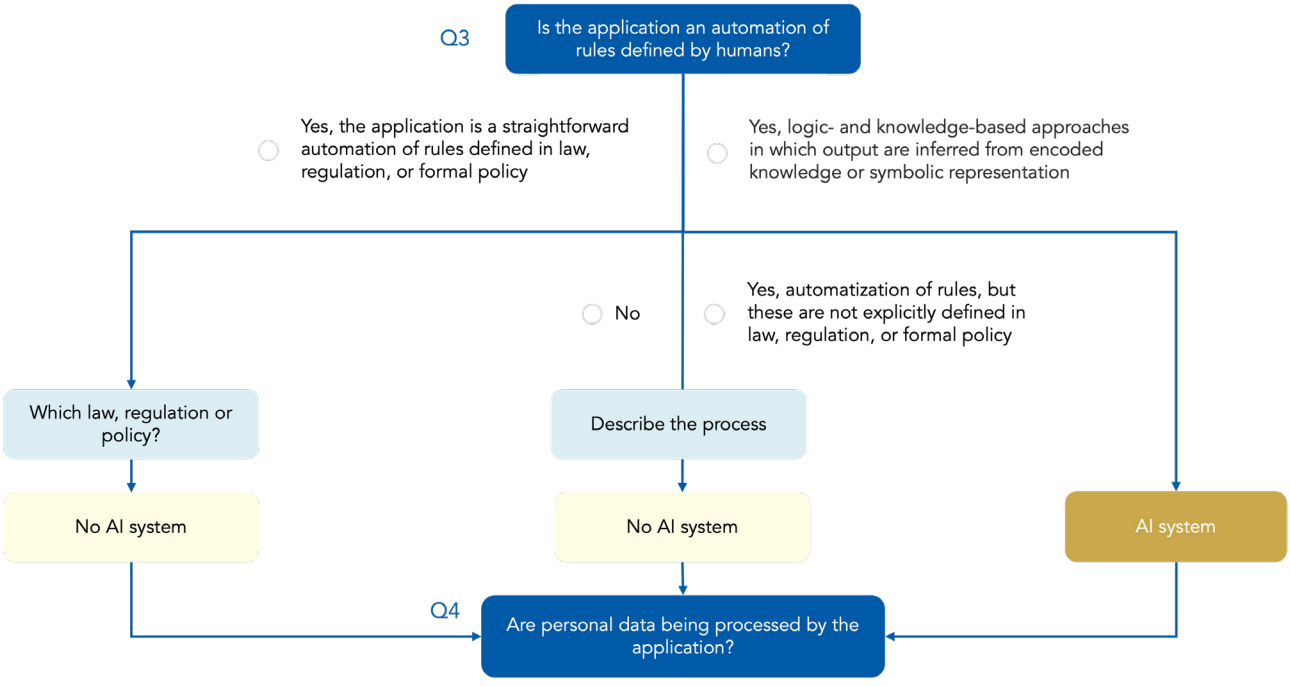

Figure 5 - Identification Q3 examines the extent to which there is human involvement in the creation of rules used in an algorithm.

conjunction with models derived from data (as already captured by Q2).[10] Purely logic- and knowledge-based approaches are rare. Developers using this type of technology are aware this represents a highly specific type. In this case the user is informed the application qualifies as an AI system. Then, users are asked whether they want to continue with the rest of Questionnaire Identification, which deals with high-impact algorithms and sADM, starting with Q8.

In case an algorithm is a straightforward automation of rules defined in law, regulation or formal policy, it does not qualify as an AI system, since there is no inference at hand. In this case, the user is first asked which official law, regulation or policy is automated, before the conclusion that no AI is involved is shared. The input provided by the user can be checked by legal experts.

In case the application is an automation of rules that are not explicitly defined in law, regulation or formal policy, the user is asked to describe how the rules were determined, before the conclusion that no AI is involved is raised again. The input provided by the user can be checked by legal experts. There is likely no inference at play. However, the distinction between rules that are explicitly stated in law, regulation or policy, and those that reflect a human interpretation or implementation of the above, is important for qualifying high-impact algorithms and sADM, as discussed in sections 3 and 4.

To assist users, the following remark is provided along Q3:

An example of rules laid down in legislation or regulations is a rule-based algorithm that, when an application for social benefits is submitted, automatically indicates whether the income and other requirements have (not) been met. In that case, the rules in the algorithm are a one-to-one implementation of norms specified in, for instance, the Dutch Participation Law.

When a standard is defined in open terms in legislation or regulations and is further specified in the application, the application does not constitute one-to-one automation of legislation or regulations.

Examples of rules defined by humans include:

> A rule-based algorithm in which a work instruction has been implemented into an algorithm;
> A risk profile in which the rules have been manually defined on the basis of employees' experience;
> Open legal standards that are further specified in rules.

Logic- and knowledge-based approaches are also referred to as 'symbolic AI systems'. This category of AI systems includes knowledge representation, inductive (logical) programming, knowledge bases, inference and deduction engines, and (symbolic) reasoning. This technology is used, for example, in expert systems.

[10] See section 3.2 of Implementation of the AI Act – Definition of an AI system, Algorithm Audit (2025).

Flowchart identification AI system

Algorithm Audit

Flowchart – AI system (Art. 3 AI Act)

This schematic representation shows the logic required to determine whether the application qualifies as an AI system according to Article 3 of the AI Act. The flowchart of the complete Identification questionnaire with all paths and outcomes can be found on the Algorithm Audit website. The complete questions can be found in the AI AQT tool itself.

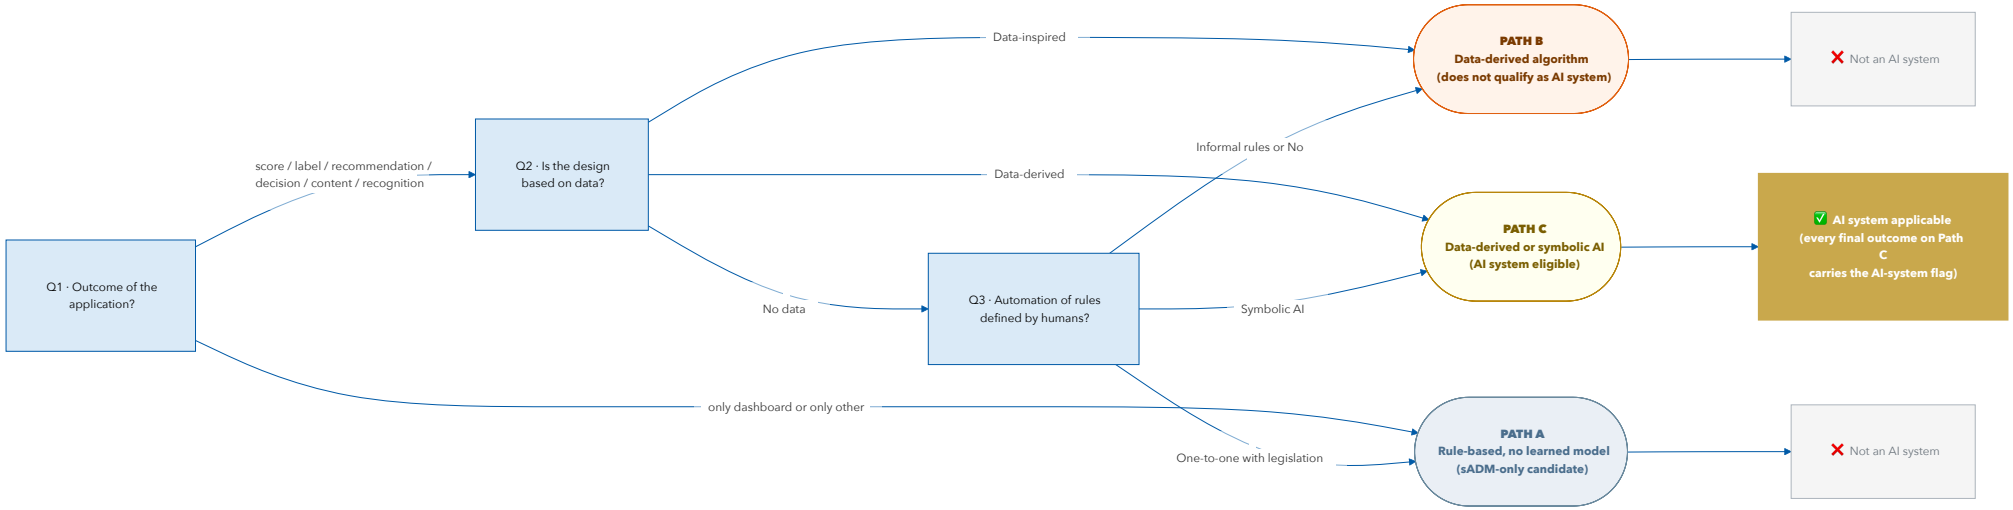

Figure 6 - Flowchart determining whether an AI system is applicable or not.

3. Questionnaire Identification – Solely automated decision-making

Article 22(1) of the General Data Protection Regulation (GDPR) states that:

"The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her."

Various resources help to clarify the scope of solely automated decision-making (sADM), including the relevant exemptions set out in article 22(2) GDPR. Competent authorities and legal experts have issued guidance how requirements can be met, such as *"Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679"*[11] issued by the European Data Protection Board (EDPB). In addition, several other sources were consulted in designing the questionnaires used in the AI AQT, including the *"Advice on Article 22 GDPR and automated selection tools"*[12] published by the Dutch Data Protection Authority (DPA), "Advice on automated selection techniques" by legal experts at Pels Rijcken[13] and relevant legal-scientific literature[14],[15]. Algorithm Audit has also published supplementary materials providing a step-by-step guide how decision based solely on automated processing can be prevented.[16]

This section first examines the core concepts that serve to identify algorithmic systems that are in scope of article 22 of the GDPR. In particular, the notion of a 'decision', 'legal effects or other significant impacts for individual' and the understanding of 'based solely on automated processing', including the latter's relationship to human intervention, are discussed (3.1). As will be shown the notions of 'decision with legal or similarly significant effect' and 'automation' from article 22 strongly overlap in meaning with those of 'direct consequences' and 'significant effect on the outcome of the process' from the Algorithm Register Guidelines for Dutch public sector organisations. As such, the same questions can be used to qualify high-impact algorithms and sADM (i.e., Q9 and Q12 in the Identification questionnaire). Hence, section 3.1. explicitly links to concepts explained in section 4. Questionnaire: High-impact algorithms.

NOTE: Not all instances of sADM qualify as a high-impact algorithm. For example, one-to-one automation can qualify as sADM but is strictly outside the scope of high-impact algorithms (see Q3 and 4.2 Significant effect on the outcome of the process).

3.1 Decision, legal or similar effects and human involvement

The scope of the prohibition in article 22 GDPR depends on at least three aspects: I. A decision is made, II. The decision has legal effects or otherwise significantly affects the individual concerned, and III. It is based solely on automated processing. Each concept is discussed in turn.

[11] Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679, European Data Protection Board (2018).

[12] Advice on Article 22 GDPR and automated selection tools, Dutch Data Protection Authority (2024).

[13] Advice on automated selection techniques, Pels Rijcken (2024).

[14] Legal protection against risk profiling based on the GDPR, the ECHR, and the Charter of Fundamental Rights, F. Çapkurt, Dutch journal for legal professionals (2025).

[15] The Right to an Explanation in Practice: Insights from Case Law for the GDPR and the AI Act, L. Metikos en J. Ausloos, Law, Innovation and Technology (2025).

[16] eaningful human intervention for risk profiling algorithms – Preventing decision-making based solely on profiling, Algorithm Audit (2025).

I. A decision is made

The notion of a decision must be interpreted broadly: not only formal decisions, as defined in Dutch public administration Law (Awb Article1:3) may affect citizens and organisations. For any type of decision-making, the individual impact must be considered (see II). The Schufa case further broadened the scope by concluding that even the computation of a score itself can constitute a decision.[17]

II. The decision has legal effects or otherwise significantly affects the individual concerned

A decision, which the output of the algorithm informs, has a 'legal effect' or otherwise 'significantly affects' individuals if one of the following types of decisions is made:[18]

- **i.** A formal decision, such as imposing a tax assessment, granting or denying a benefit or allowance, making a decision following an appeal, or granting or denying a permit or subsidy;
- **ii.** A decision with financial consequences, such as the ability to obtain a payment plan or qualify for credit;
- **iii.** Entering into an agreement, such as an employment contract or a purchase agreement;
- **iv.** Selection for an inspection, if the inspection is intrusive for the individual, such as a home visit;
- **v.** A decision affecting someone's access to education, such as admission to a university or school assignments;
- **vi.** Decisions affecting someone's employment opportunities, such as processing job applications or assigning projects to freelancers;
- **vii.** Otherwise significantly impacting the individual.

The types of decisions listed above correspond to the answer options included in Q5. It is worth noting that inspections, in particular, can have a significant effect before even concluding or without ever resulting in further actions towards an individual beyond the inspection itself. For example, an inspection can serve as grounds to delay payments or temporarily suspend access to services. A particularly intrusive inspection, like a home investigation, can have material consequences for the subject, even if not financial. Alternatively, a history of inspection can legitimatize increased caution in future cases and provoke a sequence of checks. As such, decisions to subject an individual to an inspection can be highly impactful in and of themselves and warrant detailed attention. For this reason a supplementary question, Q5.1, is raised to ascertain the secondary effects of an inspection-related decision when Q5 is answered with "Inspection, investigation, or requests for additional information".

Secondary effects might also occur if the outcomes of a risk profiling algorithm shared internally or externally, or are stored long-term significant consequences for stakeholders can arise, as discussed under the explanation for Q8-Q10 hereunder.

More information on 'legal' or 'similarly significant' effects can be found op p.21-22 of the EDPB guidelines and p.6-7 of the advice published by the Dutch DPA.

III. Solely based on automated processing

A system falls within the scope of article 22 of the GDPR when decisions are made solely through automated decision-making. One way the EDPB guidelines clarify the scope of Article 22 is by explaining on p.8 that: *"Solely automated decision-making is the ability to make decisions by technological means without human involvement"*. What human involvement entails is elaborated on p.21 of the EDPB Guidelines: *"To qualify as human involvement, the controller must ensure that any*

[17] ECLI:EU:C:2023:957, case C-634/21, Court of Justice of the European Union (2023).
[18] Supra note 12-16.

oversight of the decision is meaningful, rather than just a token gesture. It should be carried out by someone who has the authority and competence to change the decision. As part of the analysis, they should consider all the relevant data".[19]

In Q7, the concept of 'human involvement' is linked to the extent to which the algorithm influences the outcome of the process, reflecting conceptual overlap with high-impact algorithms. By distinguishing whether the algorithm's output directly or determines the process outcome (option1-2), it becomes possible to assess whether human intervention is meaningful. Further details are provided in Q7.

Although profiling, as defined in article 4(4) of the GDPR, is referenced in the legal text of article 22(1), and frequently referred to in several of the consulted resources,[20] it has no bearing on the the scope of the provision and imparts no clarifying condition for qualifying sADM.

More information on the role of work instructions for human decision makers and the relationship between automated-decision making and the obligation to prevent discrimination can be found in the paper *"Meaningful human intervention for risk profiling algorithms – Preventing decision-making based solely on profiling".*[21]

Identification Q1 – What type of output does the application derive?

The type of output generated by an algorithmic system gives an indication of whether the system qualifies as sADM. See Figure 6.

When the output is a prediction (incl. a score, ranking, label, object-, face- or voice recognition), recommendation, decision or content, these outputs are assumed to be a weighting factor in the decision-making process that utilises these outputs (see section 3.1). If one of these options is selected the user is brought to Q2.

[19] Supra note 12.
[20] Supra note 12 and 13.
[21] Supra note 18.

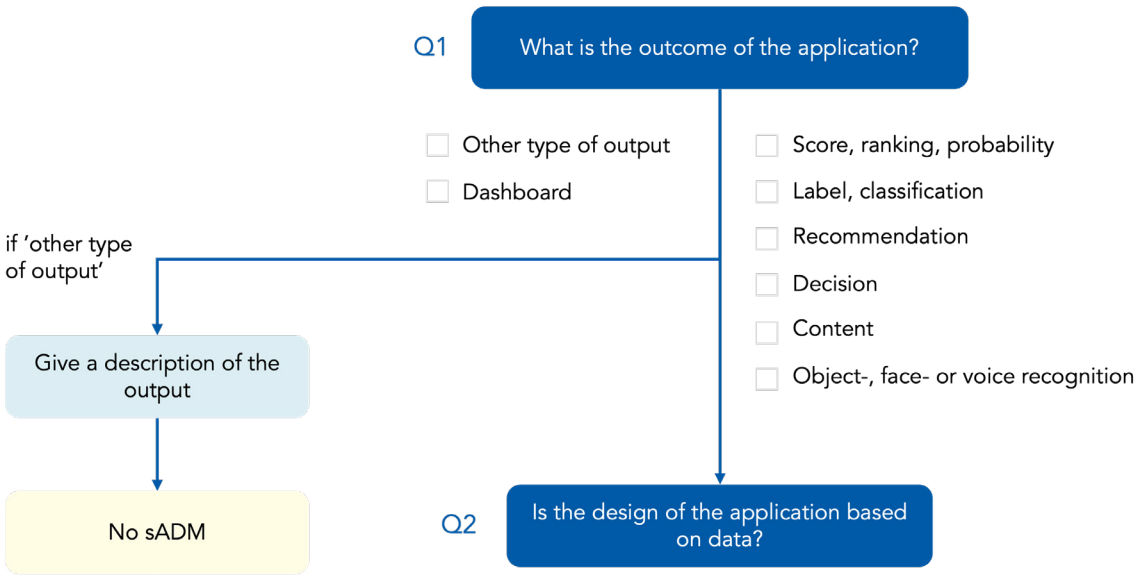

Figure 7 - Q1 of questionnaire Identification concludes that the application does not qualify as sADM if a 'Dashboard' or 'Other type of output' is selected.

Since a dashboard, on their own, only provides data visualization, there is no decision, legal or similar effects. It is the human that derives conclusions from this straightforward visualisation. If a user indicates the only output of a system is a dashboard, it is concluded it is not sADM. Users are explicitly prompted to consider if other types of outputs are displayed in this dashboard. If they select one of the specified outputs along with "Dashboard", they are brought to Q2.

The same logic applies to the option "Other type of output". If it is the sole provided answer, the tool concludes the application is not sADM. If coupled with another output (other than dashboard), the user proceeds to Q2. In either case, the user is asked to provide a description of the output, which can be manually assessed by experts.

Identification Q2 – Is the design of the application based on data?

Question Q2 provides no information for the qualification of sADM. The user answers this question because it supports qualification of AI and high-impact algorithms (see section 2 and 4).

Identification Q3 – Is the application an automation of rules defined by humans?

The only relevant difference in how Q3 captures the concepts of sADM and high-impact algorithms is that the form one-to-one automation cannot qualify as a high-impact algorithm but may still constitute sADM. More information about these cases are collected in questions Q8-Q12. In all cases users are brought to Q4 – in which information about the GDPR is collected. When the application is not an automation of rules defined by humans, users are requested to describe the process how rules are defined. See Figure 7.

Identification Q4 – Are personal data being processed by the application?

Since solely automated decision-making is a concept from the GDPR, Q4 inquiries whether personal data are being processed and whether the GDPR applies. If no personal data are processed, the process falls outside the scope of the GDPR and therefore cannot constitute sADM.

However, the absence of personal data does not mean an application is without significant impact. A neighbourhood profiling algorithm used to allocate

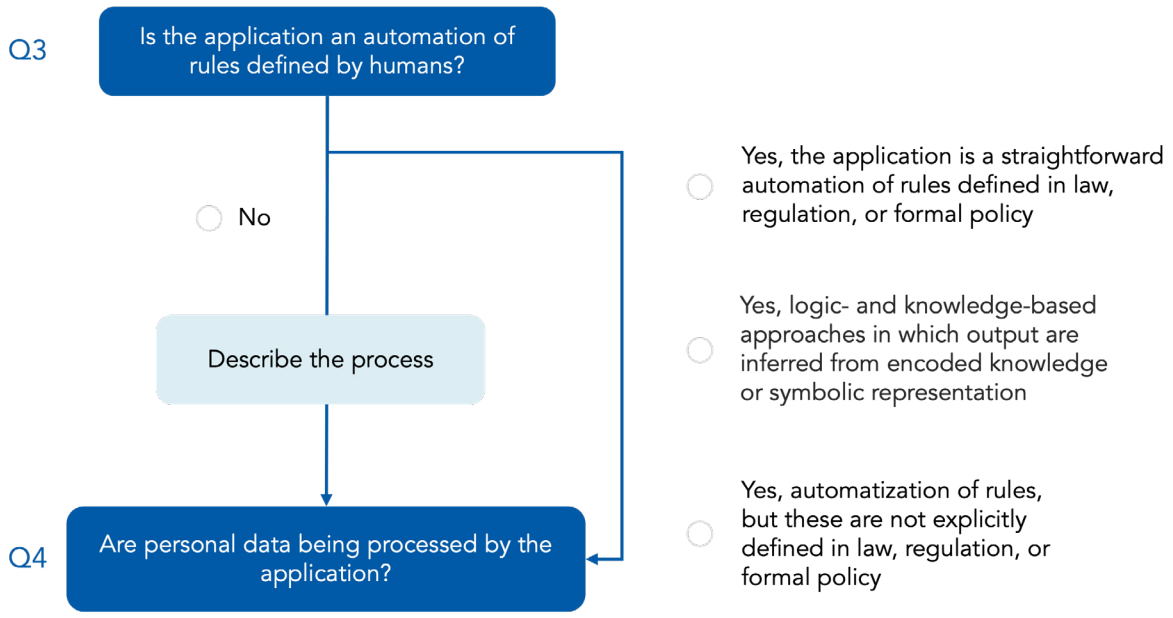

Figure 8 - All responses routes for Q3 in questionnaire Identification are directed to Q4 for GDPR identification.

police presence, for example, may have real consequences for individuals downstream. Yet, its output cannot be traced back to any specific person. Because no personal data is processed, there is no basis for prohibition under article 22 GDPR. Nevertheless, such a system may still qualify as a high-impact algorithm. These cases are discussed in section 4.

This distinction also matters in terms of scope: sADM applies only to decisions affecting individuals, whereas high-impact algorithms can affect groups as well (see section 4.1). For this reason, Q8 is included alongside Q4 to capture systems that may not trigger GDPR obligations but still carry significant societal impact. See Figure 8.

In case the output of the algorithm concerns personal data, the user is directed to Q5 to assess how the output of the algorithm is used.

If no personal data are processed, the user is still redirected to Q8. This follows the better safe than sorry principle of the design. In practice, users are expected to also answer Q8 with 'No.' Since Q8 can also capture whether a decision pertaining to an individual is made and therefore personal data should be processed. The user is not immediately presented the conclusion that sADM is not at hand when answering 'No' to Q4.

To assist users, the following remark is provided:

Examples of non-personal data are:

> Group statistics where no individual can be singled out
> Anonymised data
> Output about physical matters that are not linked to an individual, such as sensor readings, weather data, machine and operational data, urban planning and infrastructure data
> Financial records about companies
> Data about city districts and neighborhoods

NOTE: businesses in which the entrepreneur is personally liable (self-employed, sole proprietorship, general partnership, professional partnership) are considered as personal data.

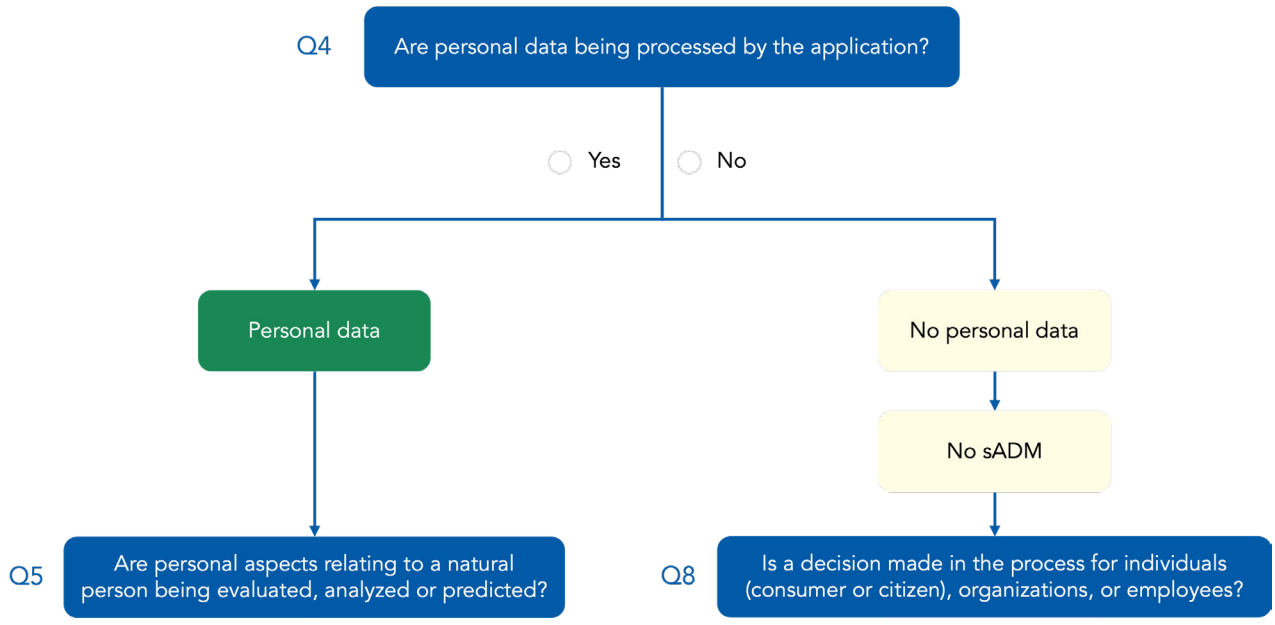

Figure 9 - Q4 of questionnaire Identification determines whether personal data are processed or not.

Identification Q5 – Are personal aspects relating to a natural person being evaluated, analyzed or predicted?

To determine whether the application involves profiling within the meaning of Article 4(4) GDPR, users are asked a question drawn directly from the legal definition. See Figure 9. In all cases, the user is forwarded to Q6.

To assist users, the following remark is provided:

Personal aspects include, for example, analysing or predicting job performance, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements.

Identification Q6 – Is the output of the algorithm shared with other organizations?

As follows from the Schufa ruling and consequent jurisprudence, sharing the output of algorithms may result in sADM (see section 3.1).[22] Q6 therefore assesses whether the output of the algorithm is shared with other organizations. See Figure 10.

When the output of the algorithm is shared with other organizations, users are presented at the end of the questionnaire a warning that sharing outputs may lead to prohibited sADM. This should be assessed with support of legal professionals.

In case the output of the algorithm is not shared with other organizations, the user is directed to Q7. <badge Identification> Q7 – Is the output of the algorithm stored for longer than the duration of the primary process for which the algorithm is used?

As follows from the Schufa ruling and consequent jurisprudence, storing the output of an algorithm may result in sADM (see section 3.1). Q7 therefore examines if the output of the algorithm is stored for longer than the duration of the primary process for which the algorithm is used. See Figure 11.

When the output of the algorithm is shared with other organizations, users are presented with a warning at the end of the questionnaire that sharing outputs may lead to prohibited sADM. This should be assessed with support of legal professionals.

In case the output of the algorithm is not shared with other organizations, the user is also directed to Q8 to continue the questionnaire.

[22] Additional explanation of the Shufa ruling can be found in supra note 17.

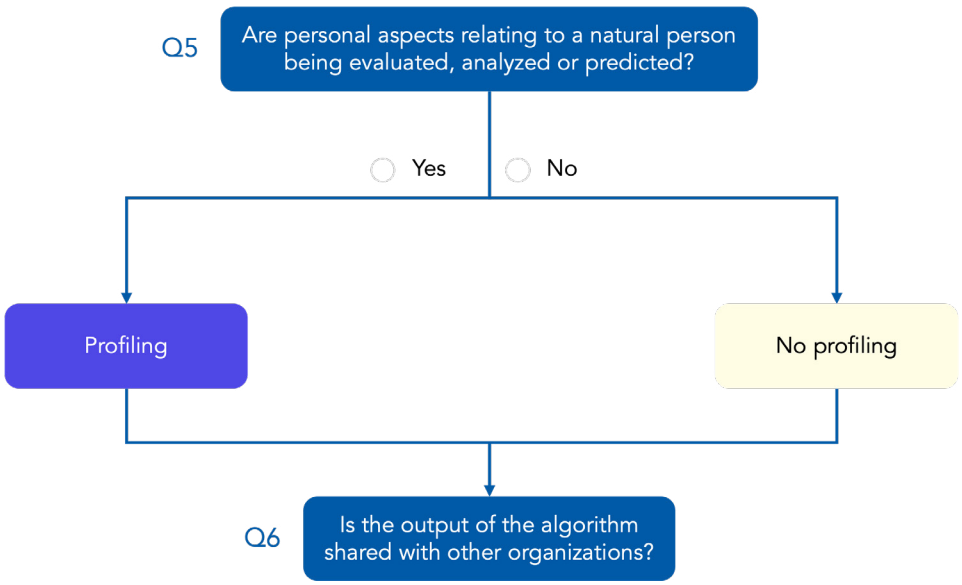

Figure 10 - Q5 of questionnaire Identification determines whether profiling, as defined in article 4(4) GDPR, is applicable.

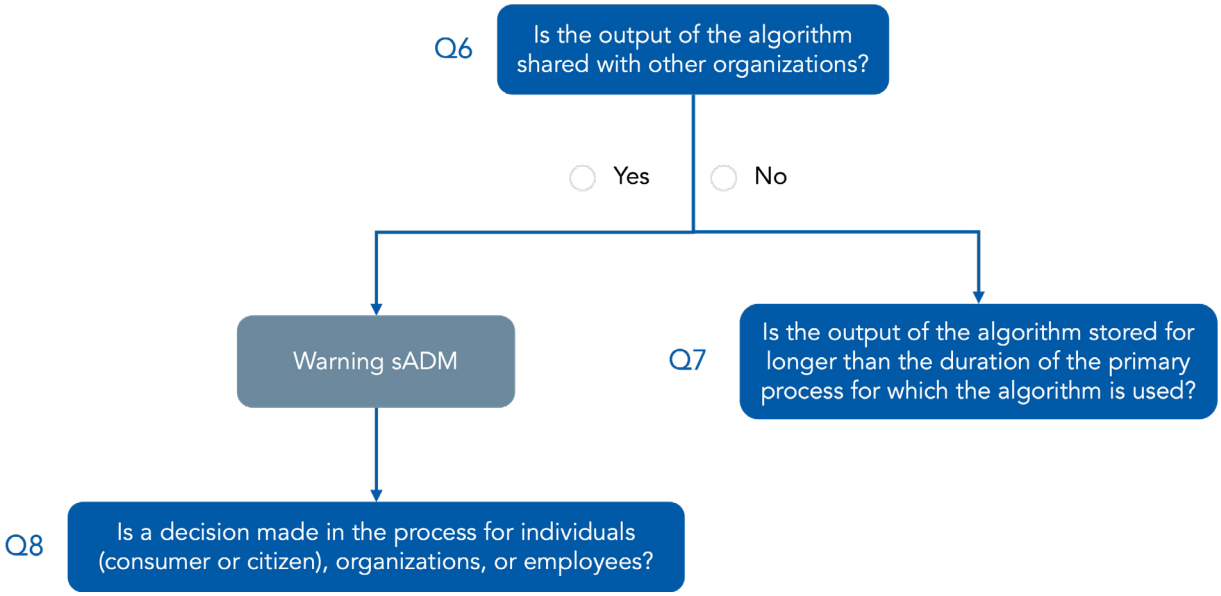

Figure 11 - Q6 of questionnaire Identification assesses whether the output of the algorithm shared with other organizations.

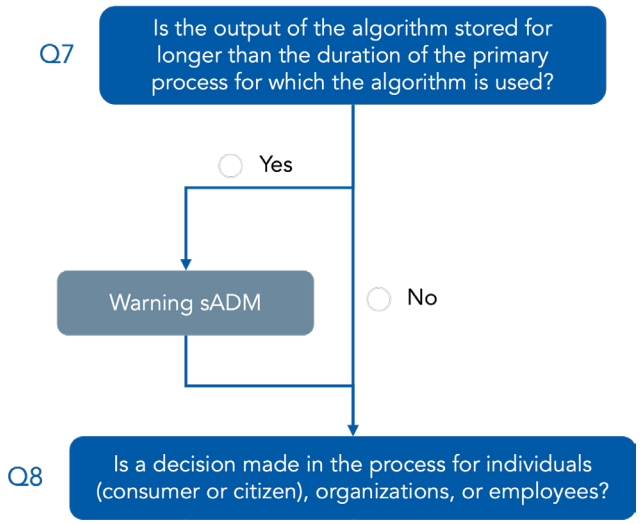

Figure 12 - Q6 of questionnaire Identification assesses whether the output of the algorithm shared with other organizations.

Identification Q8 – Is a decision made in the process for individuals (consumer or citizen), organisations or employees?

Q8 assesses whether a decision is made in the process the algorithm is involved in. This is a necessary condition for sADM to apply (see section 3.1), though not a requirement for high-impact algorithms (see section 4). If the answer is 'Yes', users proceed to Q9. If the answer is 'No', users are redirected to Q11. This routing is designed to streamline the identification of high-impact algorithms. See Figure 12.

To assist users, the following remark is provided:

For instance:

> Prioritising citizen queries or requests
> Determining whether additional information is needed from a citizen
> Selecting individuals for checks or inspections
> Assessing eligibility for services or facilities.
> For public sector organisations, the following note is also included:

NOTE: A decision is much broader than a formal decision as defined in Dutch Public Administration Law. The notion of 'decision' is also used in the context of Article 22 GDPR.

Identification Q9 – What kind of decision is made in this process?

To identify sADM, it's also needed to assess what type of decision is made in the process the algorithm is involved in. This is assessed in Q9. This serves to assess whether direct consequences follow from the decision (see section 3.1). For sADM identification, a separate answer option is included to identify secondary effect of a decision made in the process in which the application is used ('Inspection, investigation or requests for additional information'), which is further discussed in Q10. See Figure 12.

Note that 'Prioritization of applications, requests, complaints, and objections' is not considered to be sADM, because this only concerns routing of tasks which doesn't affect the decision-making process itself. This differs from the logic used for the high-impact algorithm questionnaire, as the Algorithm Registry Guidelines specifically mentions this falls in scope of the guidelines.

When 'Other or different decision' is selected, it can be concluded that no sADM occurs. The user is asked for a clarification. This can be reviewed by legal experts.

For all other answers, users are forwarded to Q12.

Identification Q10 – Is the inspection or investigation particularly intrusive for the person concerned?

As mentioned, Q10 is shown to the user only if 'Inspection, investigation, or requests for additional information' is selected in Q9. In this context, procedural decisions that do not have direct legal effects can still be considered to have a significant impact thourgh secondary effects (see section 3.1). This is assessed in Q10. See Figure 14.

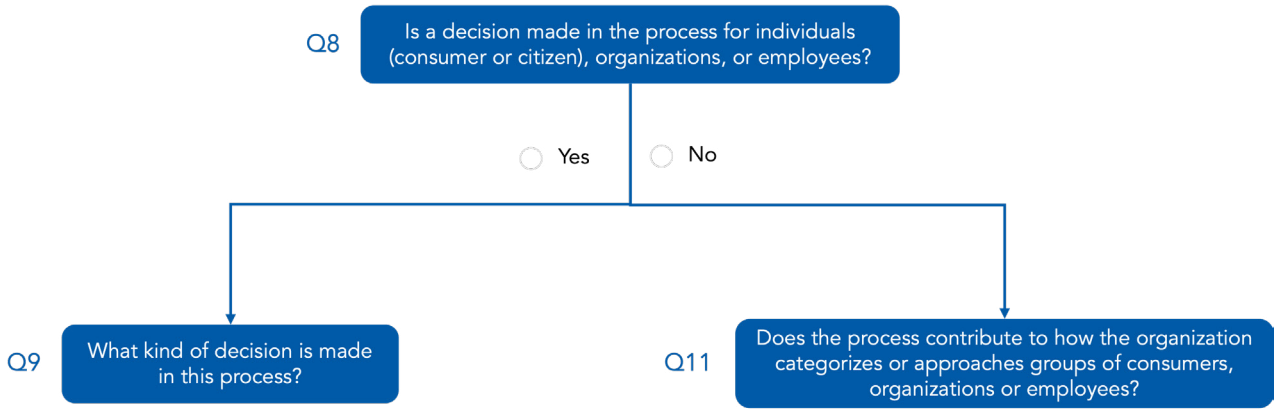

Figure 13 - Q8 of questionnaire Identification, determines whether a decision is made in the process the algorithmic application is used in.

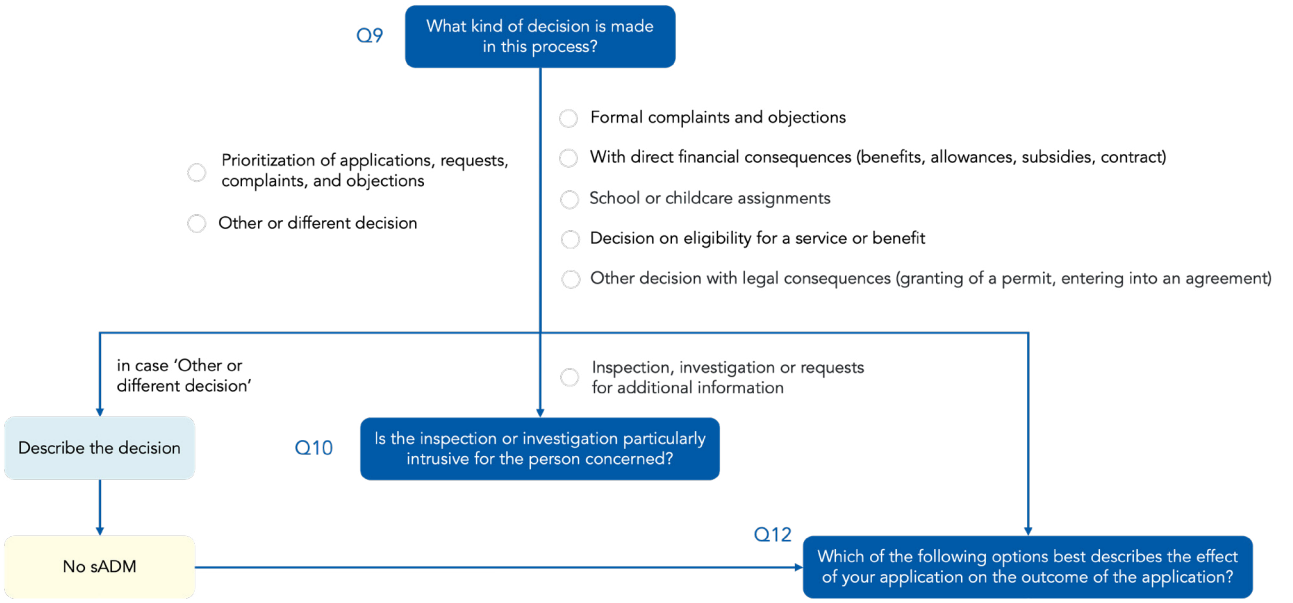

Figure 14 - ...

When being selected for an inspection or investigation results in longer waiting time or affects eligibility for (unrelated) benefits or services, the user is redirected to Q12. The same applies if the inspection or investigation makes individuals ineligible for advance payments or payment plans, or significantly affects their private life – for example, through a home visit or other highly intrusive measures.

In case the inspection or investigation is not particularly intrusive for the person concerned, sADM is not applicable. Users are still redirected to Q12 to streamline the identification of high-impact algorithms (see section 4).

Identification Q11 – Does the process contribute to how the organization categorizes or approaches groups of consumers, organizations or employees?

For public sector organisations, the question is: "Does the process contribute to how the government categorizes or approaches groups of citizens, organizations, or civil servants?"

As Q11 is only shown when Q8 is answered 'No' (no decision is made in the process for an individual (consumer or citizen), organisation or employee), sADM is not at stake. Q11 is included to streamline the identification of high-impact algorithms (see section 4).

Identification Q12 – Which of the following options best describes the effect of the application on the outcome?

In the context of sADM, Q12 assesses the role of human involvement in the decision-making process and helps determine whether sADM is at hand. See Figure 15.

In the algorithmic system directly determines or largely influences the outcome of the process, sADM is at play. In all other cases, it is concluded that sADM is not at play due to human involvement. These conclusions are shown to the user. When 'Another type of output' is selected, the user is first asked to describe the effect.

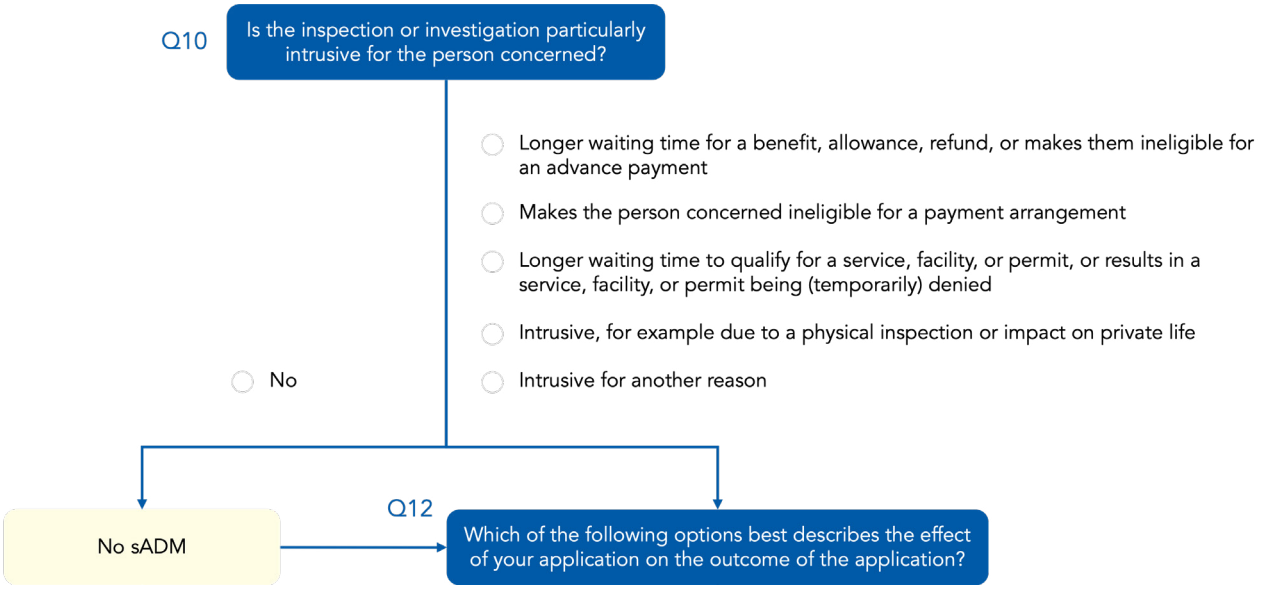

Figure 15 - Q10 of questionnaire Identification examines whether the inspection or investigation is particularly intrusive for the person concerned.

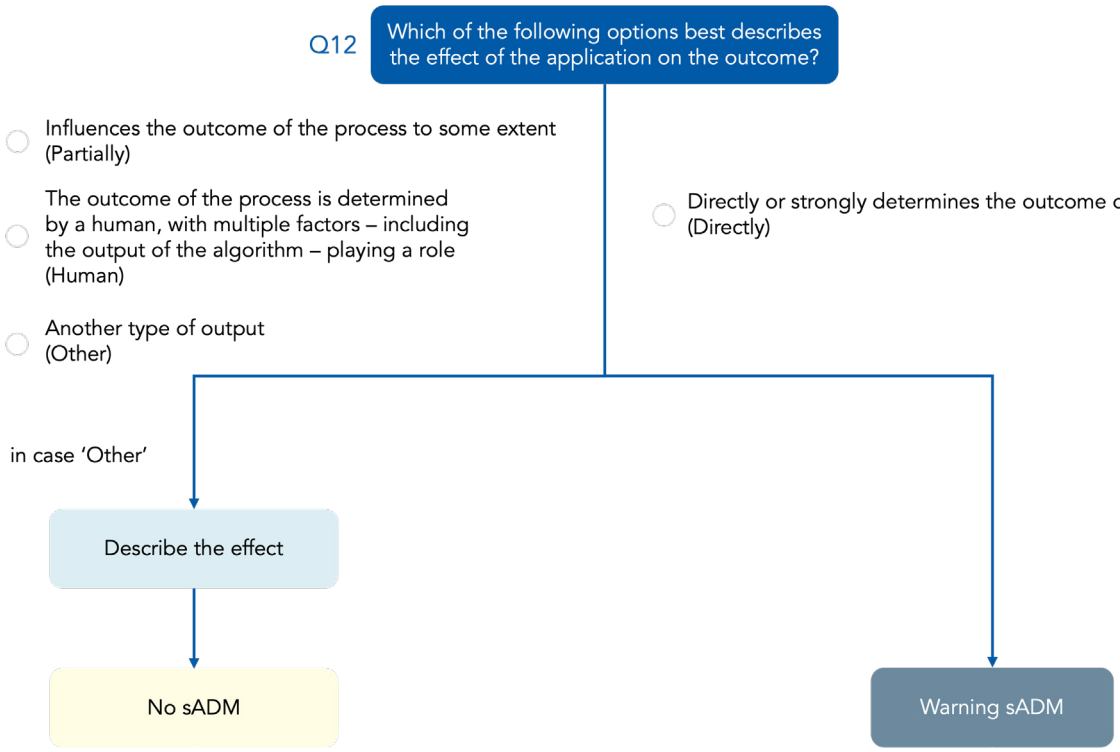

Figure 16 - Q15 of questionnaire Identification examines whether human intervention in the decision-making process is meaningful.

Flowchart identification solely automated decision-making (sADM)

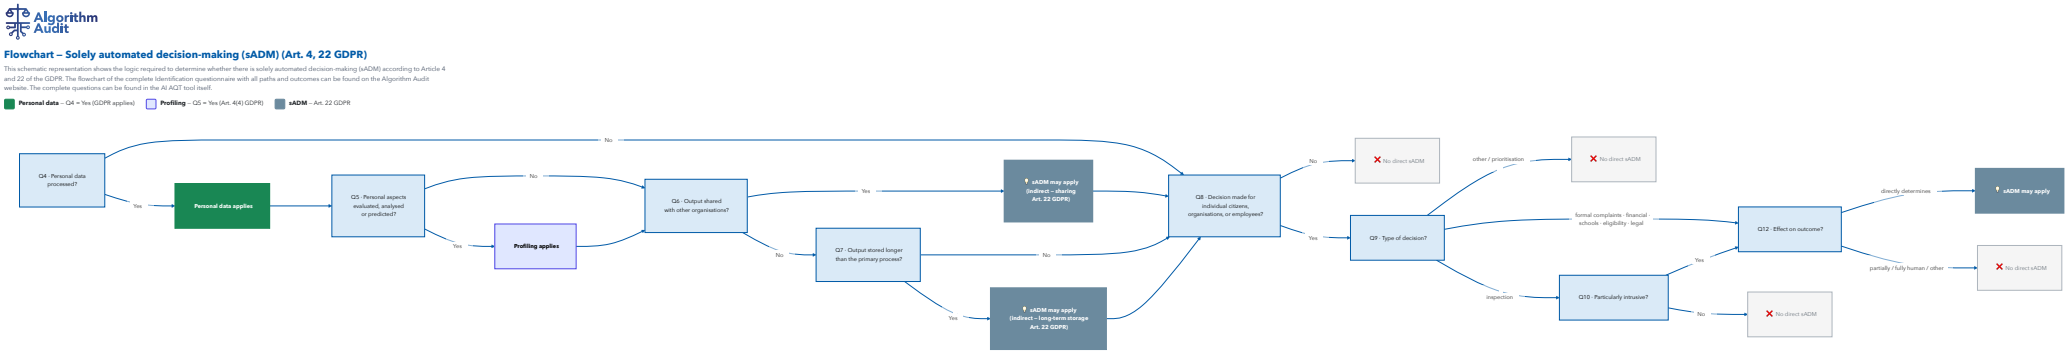

Figure 17 - Flowchart determining whether sADM is applicable or not.

4. Questionnaire Identification – High-impact algorithms

Algorithmic systems can have a significant impact on data subjects even when they do not qualify as an AI system under the AI Act (see Box 2). The term algorithm has been used since the 2010s – among others by the Dutch government – to refer to a broad category of automated systems, which also includes AI systems. In 2021, the Netherlands Court of Audit (2021) defined an algorithm as: *"A set of rules and instructions that a computer automatically follows when making calculations to solve a problem or answer a question"*.[23] An algorithmic system is a digital system through which an algorithm is executed. The definition of a high-impact algorithm, as set out in the "Algorithm Register Guidelines" issued by the Dutch Ministry of the Interior, is as follows:[24]

- **Direct consequences:** The algorithm has direct consequences for those involved (citizen, organization), e.g., imposing a fine or refusing a subsidy; or
- **Classification:** The algorithm influences how the government categorizes or approaches a data subject or group, e.g., profiling or risk indication for control.

The above categories are explained in the Algorithm Register Guidelines on the basis of three questions. See Figure 16.[25]

The three questions are the following:

1. Does it concern a process with direct consequences?
2. Are one or more algorithms used in the process?
3. Does the algorithm have a significant effect on the outcome of the process?

[23] 'Aandacht voor algoritmes', The Netherlands Court of Auditors (2021).
[24] Algorithm Register Guideline of the Ministry of the Interior and Kingdom Relations (2023).
[25] Style adapted from Algorithm Register Guideline of the Ministry of the Interior and Kingdom Relations (2023).

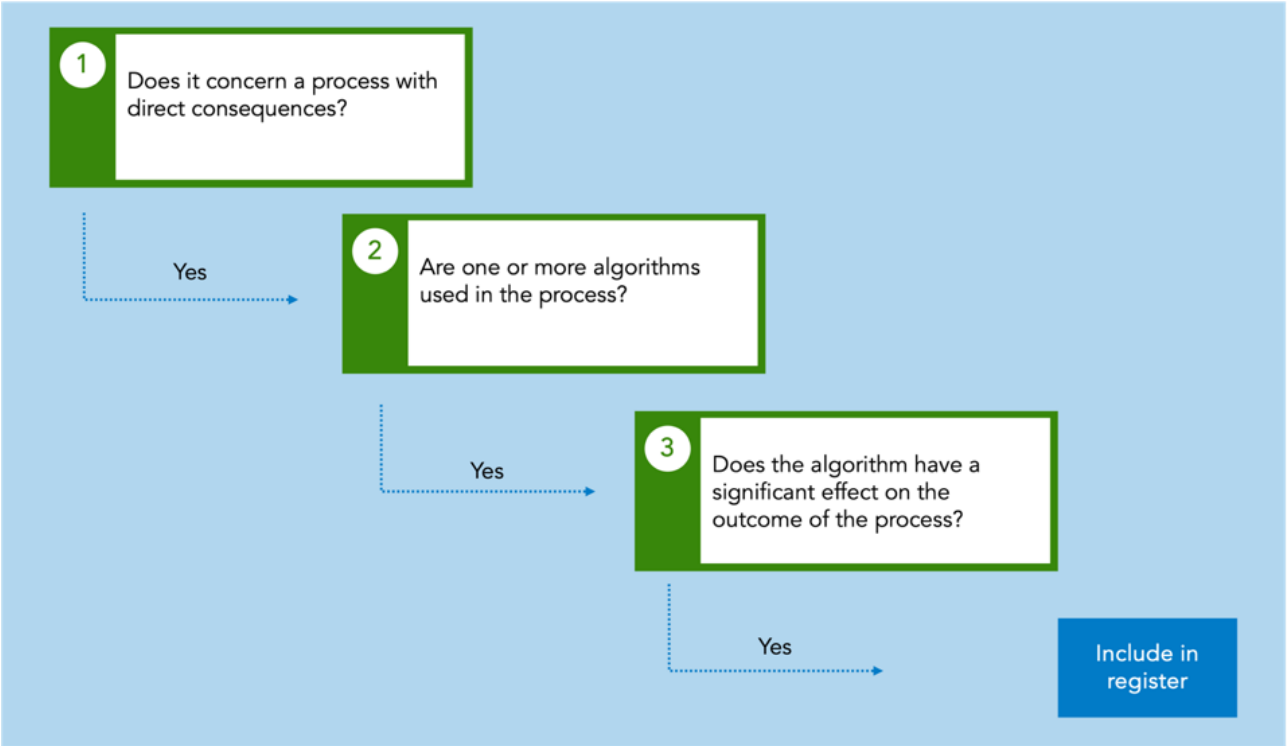

Figure 18 - Questions from Algorithm Register Guidelines that can be used to determine whether an algorithm is a 'high-impact algorithm'.

This section discusses how the first and last questions in Figure 16 – about the process in which the algorithm is used – are incorporated into the AI AQT as Questions 8-12 (Q8-12). Assuming at least one algorithm is involved in a given decision-making process, it can be deemed the second question in Figure 16 imparts no capacity to distinguish high-impact algorithms from other algorithms. The number of algorithms involved in a process is therefore not operationalised as a separate question in AI AQT. This section begins with an explanation of what is meant by the concepts of 'direct consequences' (section 4.1) and a 'significant effect' (section 4.2).

NOTE: For Dutch public sector organisations holds that high-impact algorithms must be published in the Dutch national Algorithm Register,[26] unless there is a ground for exception.[27]

[26] https://algoritmes.overheid.nl/nl/algoritme

[27] Grounds for exception as stated in the Guideline are: "legal grounds for exception as specified in the Open Government Act (Woo) and the Police Data Act (Wpg), or 'gaming the system'." See supra note 12.

Box 2

Dutch scandals were not caused by AI systems

Algorithms involved in scandals in the Netherlands, such as the childcare benefit scandal and the use of a discriminatory profiling algorithm by the Dutch Executive Agency for Education (DUO), involve algorithmic systems that fall outside the scope of the AI Act. These processes were based on human-defined rules without inference. In the case of DUO, students were assigned a risk score by a rule-based algorithm solely designed by human experts.[28] Although this algorithm does not qualify as an AI system, it had severe consequences: the risk profiling algorithm indirectly discriminated against students with a non-European migration background.[29] Likewise, the algorithm at the crux of the Dutch childcare benefit scandal also does not qualify as an AI system, yet it discriminated against Dutch citizens with dual nationality based on human defined decision rules.

High-impact algorithms appear to be more prevalent than AI systems, particularly in the Dutch public sector. In the summer of 2025, Algorithm Audit analysed the inventory of 14 Dutch ministries and concluded that 250 out of 370 (~67.6%) algorithmic systems qualify as high-impact algorithms, while only 13 out of 370 (~3.5%) qualify as high-risk AI systems.[30] This exemplifies that it is important to identify not only AI systems, but also high-impact algorithms so that appropriate control measures can be applied.

[28] Preventing prejudice, Algorithm Audit (2024).

[29] Addendum Preventing prejudice, Algorithm Audit (2024).

[30] Inventory 14 Dutch Ministries Netherlands Algorithm Registry, Algorithm Audit (2025).

4.1 Direct consequences

A high-impact algorithm is used in a process that has direct consequences for those involved. As clarified in the Algorithm Register Guidelines:[31]

I. *"These are processes with impact, which will generally be decision-making processes. Or the process contributes to how the government categorizes or approaches a person or group, for example by using weighting factors or predictions. This can have consequences for the approach or treatment. Examples of the latter are risk assessments and algorithms for fraud detection."*

II. *"In any case, the consequences include legal consequences. A legal consequence means that the decision under the Dutch Public Administration Law (Algemene wet bestuursrecht) affects the legal rights of a data subject, a person's legal status or their rights under an agreement. It also concerns factual consequences that affect the interests of a person, such as financial consequences (whether or not to receive an allowance), consequences for fundamental rights (whether or not to provide legal protection) and legal consequences (whether or not to stay in the Netherlands, to be allocated a home). The selection for an inspection or control is also seen as a consequence."*

III. *"Stakeholders include everyone who has to deal with the Dutch government. We summarize this as citizens and organizations."*

Direct consequences are broadly defined, which means many systems may qualify as high-impact algorithms. To determine whether an algorithm-driven process involves direct consequences, the first step is to assess whether a decision is made (Q8). For Dutch public sector organisations, the notion of a decision here is to be interpreted broadly: not only formal decisions, as defined in Dutch public administration Law (Awb Article1:3), may affect citizens and organisations. Other types of decisions may also have significant consequences for those involved and can therefore indicate the presence of a high-impact algorithm.

Once it has been established whether the process involves a decision affecting individual citizens or civil servants, the nature of the decision must be examined. This may include, for example, prioritisation of cases, decisions on formal complaints or objections or decisions with financial consequences (Q9).

Even in cases where no decision results from the algorithm-driven process, the system may constitute a high-impact algorithm if it alters how the government handles or approaches data subjects. Whether such an effect is present is assessed in Question 11 (Q11).

The following section discusses how the extent to which the algorithm impacts the outcome of the process is used to distinguishing high-impact algorithms from other algorithms.

4.2 Significant effect on the outcome of the process

A high-impact algorithm has a significant effect on the outcome of the process. As clarified in the "Algorithm Register Guidelines":

I. *"This is not about processes in which the algorithm automates/digitizes a manual work instruction. Such as algorithms in which all parameters are legally fixed and the algorithm runs through a (complex) decision tree based solely on these parameters".*

II. *"This does concern processes in which the algorithm influences a decision. Such as algorithms in which a weighting factor is given that (partly) determines the next step in the*

[31] Supra note 15.

process. The weighting factors are filled in by the space or freedom that an administrative body is entitled to in carrying out its tasks."

As explained in I., automation of rules that are 'legally fixed' in law, regulation or formal policy falls outside the scope of a high-impact algorithm. This can be dubbed 'one-to-one automation.' In such cases, the algorithm has no influence on the process – performed manually, the outcome would be no different. However, there are cases of '0.8-to-1 automation', where there is room for an organization to interpret a provision itself and formulate it as a decision rule. Such systems may qualify as a high-impact algorithm since any arbitrary specifications embedded into the algorithm shape the process. How a rule-based algorithm is design is assessed in Q3. Additional information about the effect of the algorithm on the outcome of the process is derived from Q12.

NOTE: The logic for the flowchart for the answer options in Q3 are different when assessing high-impact algorithms than when assessing an AI system. See 2. Questionnaire: AI system.

Identification Q1 – What type of output does the application derive?

The type of output generated by an algorithmic system gives an indication of whether the system influences the outcome of a process. Therefore, Q1 is used to determine whether a system could be considered a high-impact algorithm. See Figure 17.

When the output is a prediction (incl. a score, ranking, label, object-, face- or voice recognition), recommendation, decision or content, these outputs are assumed to be a weighting factor in the process that utilises these outputs (see section 4.2.II). If one of these options is selected the user is brought to Q2.

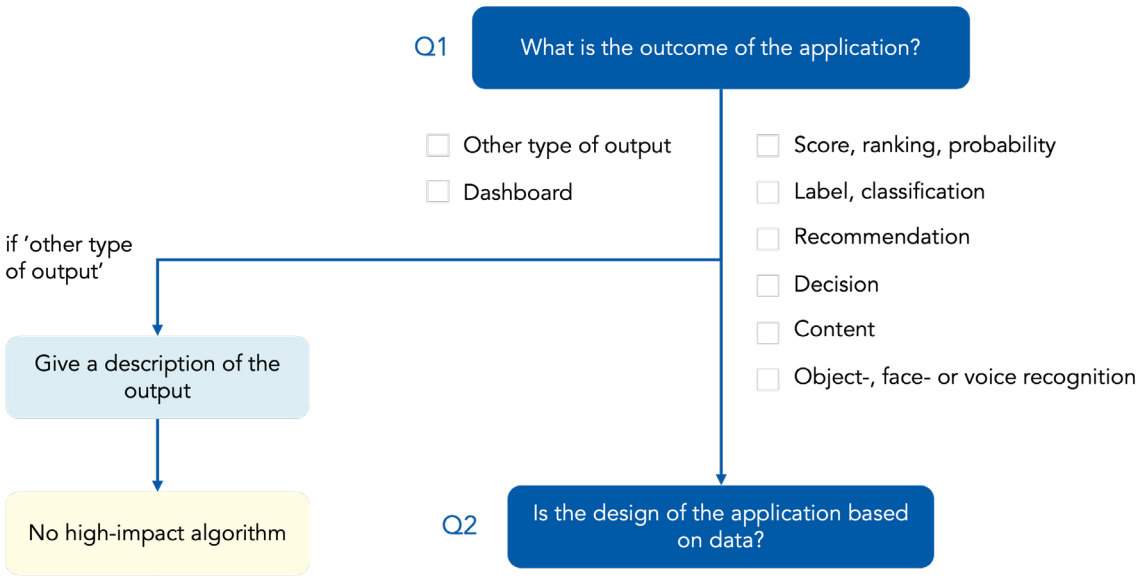

Figure 19 - Q1 concludes that the application does not qualify as sADM if a 'Dashboard' or 'Other type of output' is selected.

Since a dashboard, on their own, only provides data visualization, there is no direct consequence or significant effect at stake. It is the human that derives conclusions from this straightforward visualisation. If a user indicates the only output of a system is a dashboard, it is concluded it is not a high-impact algorithm. Users are explicitly prompted to consider if other types of outputs are displayed in this dashboard. If they select one of the specified outputs along with "Dashboard", they are brought to Q2.

The same logic applies to the option "Other type of output". If it is the sole provided answer, the tool concludes the application is not a high-impact algorithm. If coupled with another output (other than dashboard), the user proceeds to Q2. In either case, the user is asked to provide a description of the output, which can be manually assessed by experts.

Identification Q2 – Is the design of the application based on data?

If the algorithmic system is derived from data, it is not a straightforward automation of policy (see section 4.2.I). Therefore, Q2 supports qualification of high-impact algorithms.

If a system is derived from data or if the application contains components derived from data, the user is prompted to provide additional clarification and is sent to Q4.

When a user picks the option "no", they are sent to Q3 to understand whether this system should be excluded from the qualification of high-impact algorithm.

Identification Q3 – Is the application an automation of human-defined rules?

Whether the algorithmic system automates rules specified in law, regulation or formal policy or whether rules are defined by policy makers is an important indicator whether a system qualifies as a high-impact algorithm. See section 4.2 Significant effect on the outcome of the process and Figure 19 In case the application is a straightforward automation of rules defined in law, regulation, or formal policy, the user must specify in which policy instrument the rule is stated. Thereafter, it is concluded that no high-impact algorithm is at stake (see section 4.2 Significant effect on the outcome of the process).

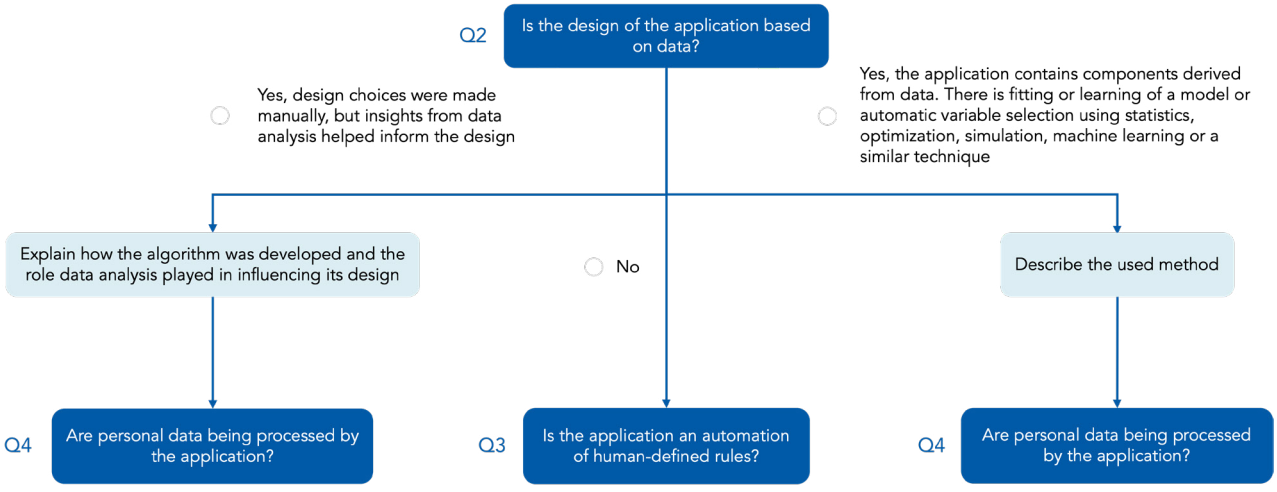

Figure 20 - Q2 of questionnaire Identification distinguishes whether more information is needed about the design of the algorithm (Q3), or that the user can continue to the next part of the questionnaire (Q4).

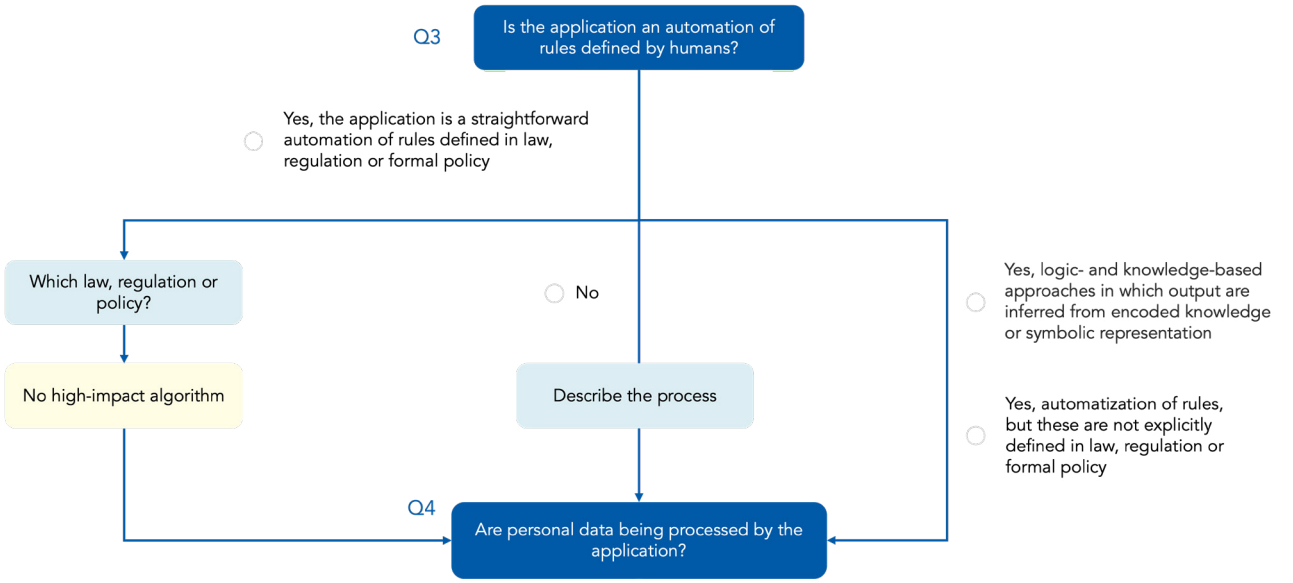

Figure 21 - In case of 1-to-1 automation of defined rules in law, regulation or formal policy, Q3 of questionnaire Identification concludes that the application is no high-impact algorithm. Other scenarios are forwarded to Q4.

If rules are not explicitly defined in law, regulation or formal policy, the user is forwarded straight away to Q4.

In case the application is not automation of human-defined rules, the user must describe the process and is then redirected to Q4.

An explanation for the answer option covering logic- and knowledge-based approaches can be found in 2. Questionnaire Identification: AI system, as this option is only relevant for AI systems.

NOTE: For this question, AI AQT recommends that work instructions for human decision-makers fall under answer option 'automatization of rules, but these are not explicitly defined in law, regulation or formal policy' rather than 'straightforward automation of rules defined in law, regulation or formal policy', as suggested in the Algorithm Registry Guidelines.

Identification Q4-Q7

Q4–Q7 don't play a role in identifying a high-impact algorithm. These questions are relevant for identifying sADM. A description of these questions can be found in section 3. Questionnaire: Solely automated decision-making.

Identification Q8 – Is in the process a decision made for individual citizens, organizations or employees?

To identify an essential characteristic of a high-impact algorithm (see section 4.1), it is first assessed whether a decision is made in the process the algorithm is involved. See Figure 20.

In case a decision is made, the user is forwarded to Q9. In case a decision is not made, the user is forwarded to Q11.

User guidance for this question is discussed in section 3. Q8.

Q4 relevant for qualifying high-impact and sADM

Q4 also supports the qualification of an algorithmic as sADM. This is further explained in 3. Questionnaire Identification – Solely automated decision-making.

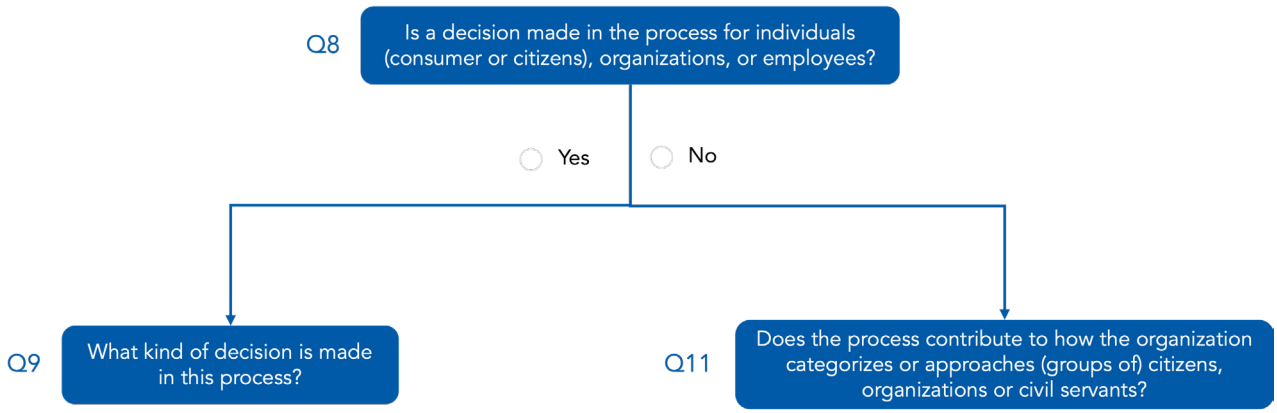

Figure 22 - Q8 of the questionnaire Identification examines whether a decision is made in the process in which an algorithm is involved.

Identification Q9 – What kind of decision is made in this process?

When a decision is made in the process the algorithm is involved in, it is important to assess the type of decision. See section 4.1 Direct consequencess and Figure 21.

A list of options is provided to help users identify the type of decision that can result in direct consequences for the data subject and may therefore indicate a high-impact algorithm. This list of options has been designed in collaboration with the municipality of Amsterdam, reflecting the breadth of decision-making with significant impact in the public sector. If one of these options is selected, the user is directed to Q12. Note that only one option can be selected. In future versions of AI AQT, answer options can be tailored to the specific sector in which the tool is used.

If a type of decision does not fall within any of the listed categories, it is concluded that the potential for direct consequences onto the subject is limited and that the algorithm in question does not qualify as a high-impact algorithm. In such cases, the user is requested to describe the type of decision, and the conclusion is shared with the user.

NOTE: Q9 also supports the qualification of a system as sADM. This is further explained in 3. Questionnaire: Solely automated decision-making.

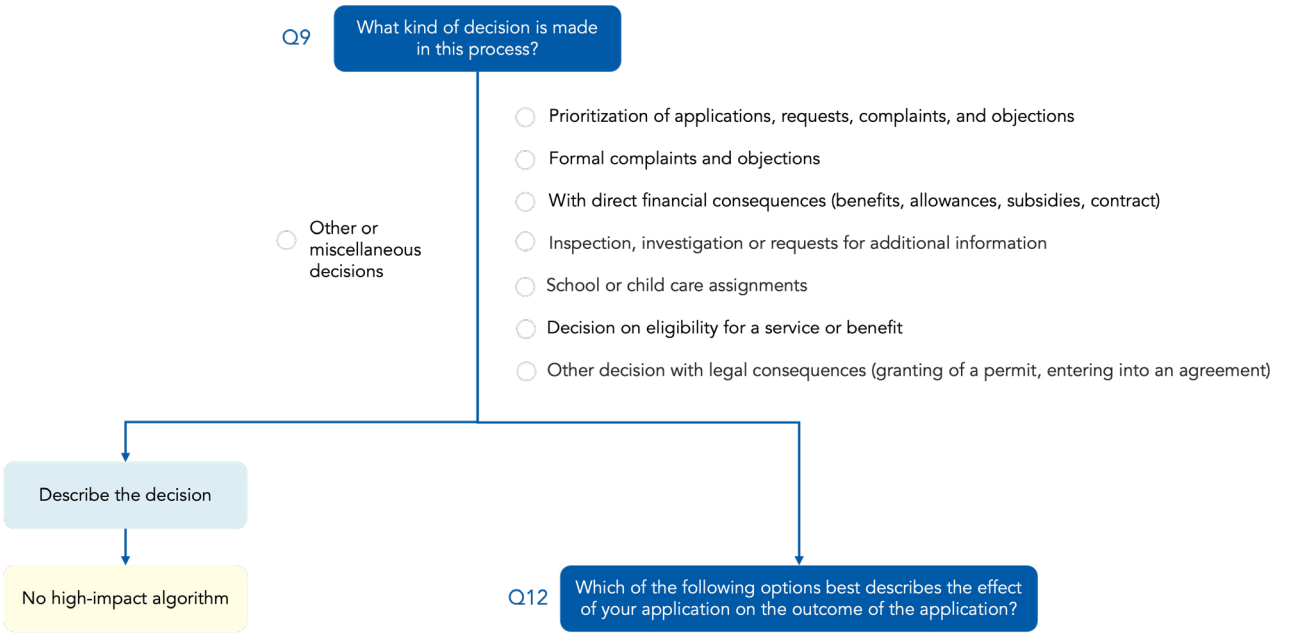

Figure 23 - Q5 of questionnaire Identification asks what kind of decision is made in the process in which the algorithm is involved.

Identification Q10 – Is the inspection or investigation particularly intrusive for the person concerned?

Q10 doesn't play a role in identifying a high-impact algorithm. This question is relevant for identifying sADM. A description of this question can be found in section 3. Questionnaire: Solely automated decision-making.

Identification Q11 – Does the process contribute to how the organization categorizes or approaches groups of consumers, organizations or employees?

For the public sector profile, the question is: "Does the process contribute to how the government categorizes or approaches (groups of) citizens or civil servants?".

Given that no decision is made in the process for an individual citizen, organisation or employee, as results from Q8, Q11 examines whether the process contributes to how the organization categorises or approaches groups of individuals, organisations or civil servants. See Figure 22. This is the second aspect of checking whether a direct consequence might follow from applying the algorithm (see section 4.1).

If the process does not contribute to how the government categorises or approaches groups of consumers, organisations or employees, it is concluded that the algorithm in question is not a high-impact algorithm and this conclusion is shared with the user.

If this cannot be said with certainty, an explanation is requested, after which the user is directed to Q12. The user will also be redirected to Q12 if Q11 is answered with 'Yes'.

NOTE: Q11 deliberately has a broader scope than individuals alone. For example, profiling neighborhoods to assign police surveillance capacity does not result in a decision about individuals, but it does influence how groups of citizens or consumers are approached by organizations or governmental actors.

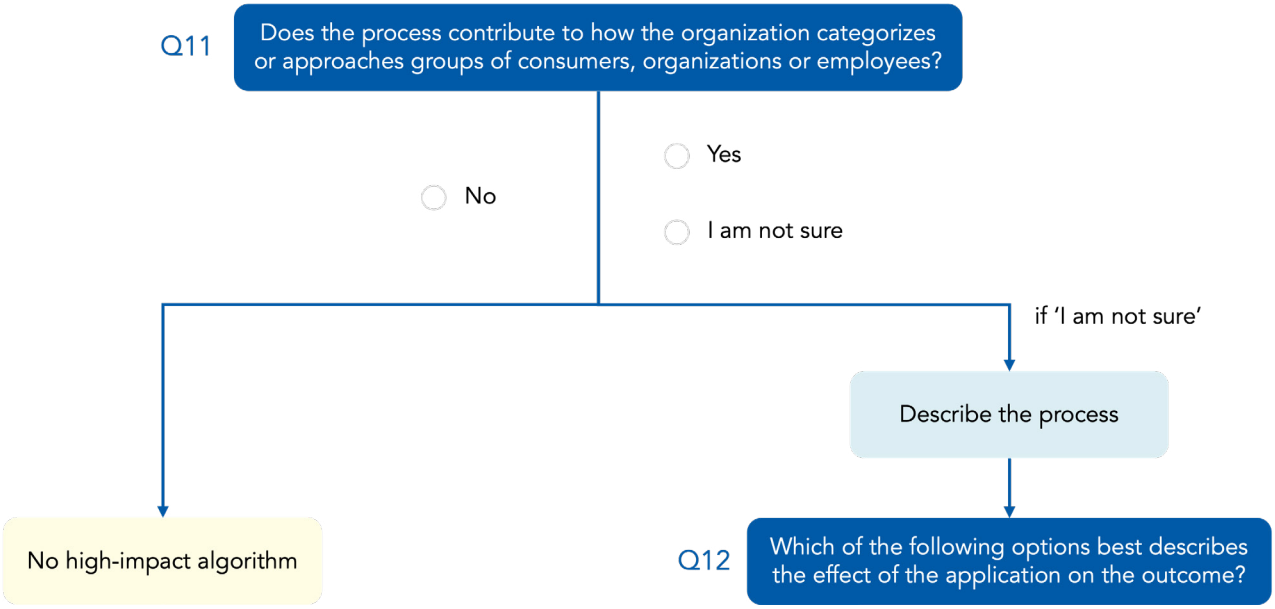

Figure 24 - Q11 of questionnaire Identification examines whether the process contributes to how groups of citizens, organisations or civil servants are categorised or approached by the government.

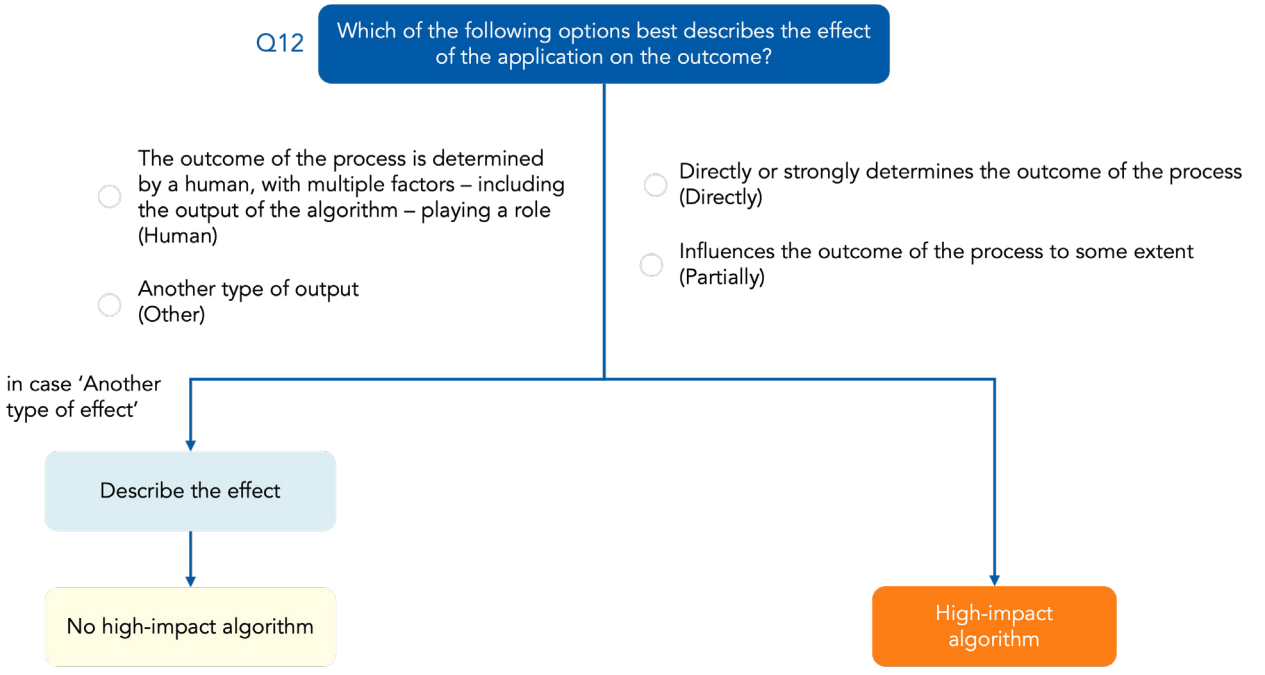

Figure 25 - Q12 of questionnaire Identification examines what the effect of the outcome of the algorithm is on the outcome of the process.

Identification **Q12 – Which of the following options best describes the effect of the application on the outcome?**

After it is established whether a decision is taken (Q8 and Q9) or the process alters how the organization approaches data subject (Q11),it must be determined whether the algorithm significantly effects the outcome of the process. If both conditions are met, it can be concluded that the algorithm is a high-impact algorithm. See Figure 23.

The scenarios described in the answer options assist users in selecting the most relevant description of the algorithm's effect on the outcome of the process. Where the outcome of the algorithm directly determines or partially influences the outcome to some extent, it is considered to have a significant effect and therefore qualifies as a high-impact algorithm.

Where the outcome of the process is determined by a human and multiple factors including the algorithm's output – play a role, the application is considered to not be a high-impact algorithm. The same conclusion applies where the application generates a type of effect other than those listed in the options. In such cases, users are asked to describe the effect so that the situation can be assessed manually by an expert team.

NOTE: Q12 also supports the qualification of a system as sADM. This is further explained in 3. Questionnaire: Solely automated decision-making.

Flowchart identification high-impact algorithm

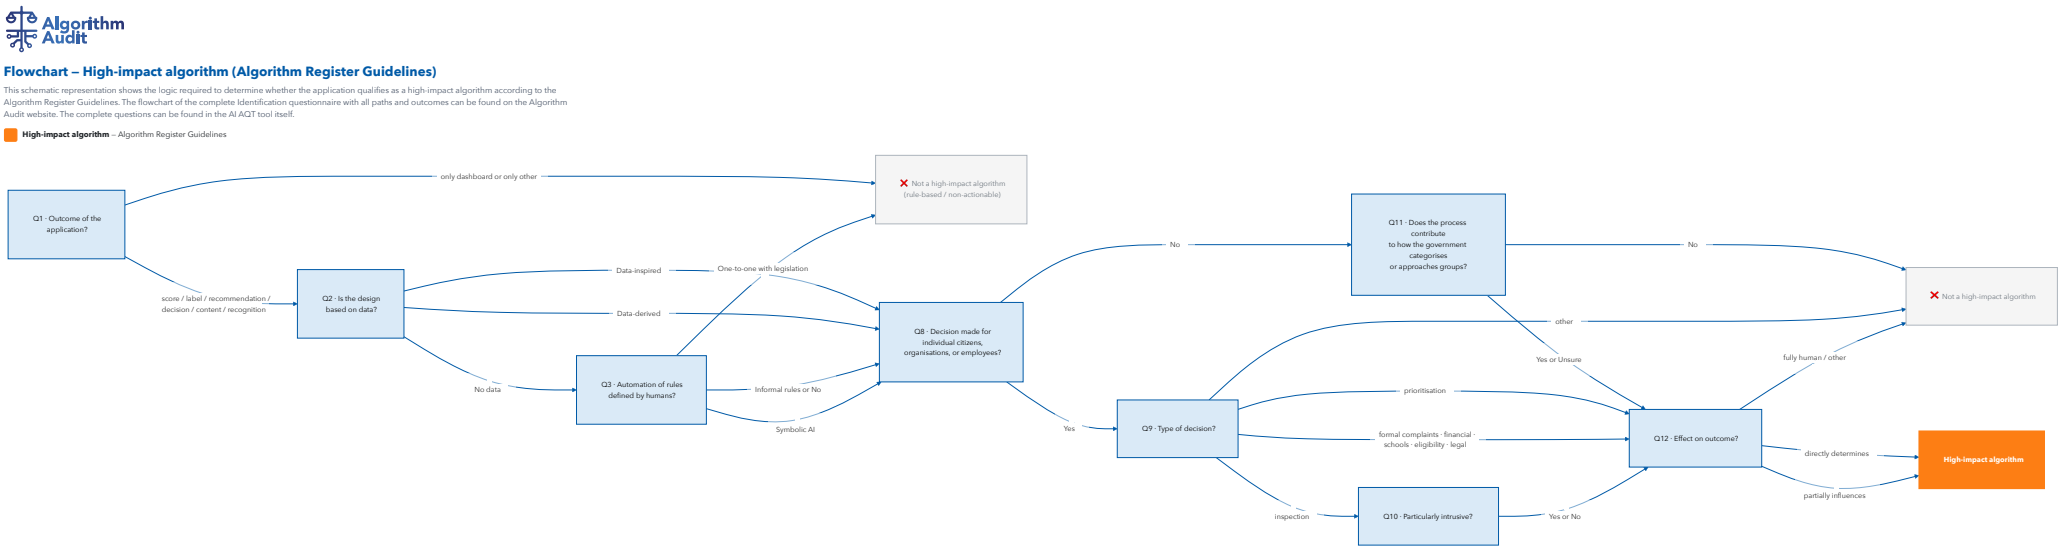

Figure 26 - Flowchart determining whether a high-impact algorithm is applicable or not.

5. Questionnaire
Role and status

The AI Act assigns obligations to specific actors involved in the lifecycle of an AI system. Articles 2 (Scope) and 3 (Definitions) of the AI Act describe these actors and the conditions under which the regulation applies to them. Once an algorithmic application has been identified as an AI system in Questionnaire Identification, the next step is to determine which role is held in relation to that AI system, since the role determines which obligations apply (Articles 16, 23, 24 and 26 AI Act).

The AI Act distinguishes five operator roles (Article 3 AI Act): provider, deployer, authorised representative, importer and distributor. A natural person who uses an AI system in the course of a purely personal, non-professional activity falls outside the scope of the AI Act (Article 2(10) AI Act). In addition to the role, the status of the AI system (whether it is already in use or still in development) determines the timelines for compliance (Articles 111 and 113 AI Act). AI systems that were already in use before the AI Act entered into force are subject to the transitional provisions of Article 111 and must comply with the requirements of the AI Act by 2030 at the latest. New AI systems face shorter, type-dependent deadlines under Article 113. Most notably, the deadline for most high-risk applications is 2 August 2027, while the prohibition on banned AI practices has applied since 2 February 2025.

This section first explains the operator roles (section 5.1) and the different usage status of an AI system (section 5.2). It then describes how these are operationalised through three questions (Q1, Q1.1 and Q2), which together make up the Role and status questionnaire.

5.1 Operator roles

It's the goal of questionnaire Role and status to assign one of the following operator roles, as stated in Article 3 of the AI Act:

- **Provider** – Natural or legal person that develops an AI system or general-purpose AI model and places it on the market under its own name or trademark (Article 3(3) AI Act).
- **Deployer** – Natural or legal person using an AI system under its own authority in a professional context (Article 3(4) AI Act).
- **Authorized representative** – A natural or legal person located or established in the Union who has received and accepted a written mandate from a provider of an AI system to carry out on its behalf the obligations and procedures established by the AI Act (Article 3(5) AI Act).
- **Importer** – Natural or legal person established in the Union who places on the market an AI system bearing the name or trademark of a non-EU entity (Article 3(6) AI Act).
- **Distributor** – Natural or legal person in the supply chain that makes an AI system available on the Union market without altering its properties (Article 3(7) AI Act).
- **Private user** – AI system is used in the course of a personal, non-professional activity (Article 3(4) AI Act).

Multiple roles can apply simultaneously. For example, an organisation that develops an AI system and uses it itself is at the same time provider and deployer. As Article 3(11) explains that "putting into service" is defined as the supply of an AI system for first use directly to the deployer or for own use in the Union for its intended purpose. So, an organisation that develops an AI system and starts using it internally is "putting it into service", which makes it a provider.

In this questionnaire, a 'Product manufacturer' is treated as a 'Provider'. A product manufacturer places an AI system on the market or puts it into service together with its product and under its own name or trademark is treated as a provider (Article 25(3)). Whether this exception applies depends on whether the AI system relates to products covered by EU harmonisation legislation listed in Annex I of the AI Act. Whether this is the case is determined in the questionnaire Risk category.

In this questionnaire, a 'Downstream provider' is also treated as a 'Provider'. A downstream provider is a provider that integrates a general-purpose AI (GPAI) model from a third party into its own AI system (Article 3(68) AI Act). Whether a GPAI model, in the form of a generative or interactive AI, is involved is addressed in questionnaire Risk category.

5.2 Status of the AI system

The AI Act draws a distinction between AI systems that are already in use and those that are still in development at the moment its obligations enter into force. This distinction is anchored in two provisions:

- **In use** – High-risk AI systems already placed on the market or put into service before 2 August 2027 must comply with the requirements of the AI Act by 2 August 2030 at the latest, unless the system is significantly changed in its design after that date (Article 111). Whether the high-risk classification applies is determined in questionnaire Risk category.
- **In development** – Read together with Recitals 177-179, sets out requirements for new AI systems, with type-dependent compliance deadlines (Article 113).

"In use" means that an AI system is actively being applied, including in a pilot or testing setting that affects work processes. Otherwise, the system is considered "in development".

Role and status **Q1 – How is the AI system developed or deployed?**

Q1 asks how the AI system is developed or deployed. The answer to this question is the primary determinant of the operator role. See Figure 27.

The answer options are mapped to operator roles as follows. Seven answer options assign one or two roles directly:

- "We only develop the AI system" → Provider.
- "We develop and use the AI system ourselves" → Provider and Deployer.
- "We develop the AI system and others also use it" → Provider and Deployer.
- "We import the AI system from a supplier outside the EU" → Importer
- "We buy the AI system from a supplier within the EU and resell it" → Distributor
- "We represent a non-EU person or entity to carry out AI Act obligations." → Authorised representative
- "You are a natural person using an AI system for purely personal non-professional activity." → Private user.

For all of the above answer options, after having answered Q1, users are redirected to Q3 to determine the status of the AI system.

The remaining answer option ("We use an externally developed AI system") routes to a follow-up question about responsibilities along the AI value chain (Q2). This question needs to be answered before an operator role is assigned.

Role and status **Q2 – Does one of the following scenarios apply to your AI system?**

Q2 determines if situations are applicable that a distributor, importer, deployer is considered to be a provider. This is the case when the following scenarios apply:

- Name and trademark: The name or trademark of your organization is put on the AI system (Article 25(1)(a)).
- **Substantial modification:** A substantial modification has been made to the AI system (Article 25(1)(b)). Beyond the formal legal definition in Article 3(23), the AI AQT describes a of a substantial modification as a significant change to an AI system's architecture, training data or functionality that materially alters how the system performs or the outputs it produces. When a high-risk AI system continues to learn after being placed on the market or put into service, this is not considered a substantial modification, provided the changes were pre-determined at the moment of the initial conformity assessment and described in the initial technical documentation (Article 43(4) AI Act).
- **Intended purpose:** The intended purpose of the AI system is modified by your organization (Article 25(1)(c)).

If one of the above scenarios apply, 'Deployer' is being assigned as the operator role. If answer option 'None of the above' is selected, the 'Provider' operator role is assigned. See Figure 27.

Whether a distributor, importer or deployer becomes a provider under the above scenarios mainly matters for high-risk AI systems. The Risk category questionnaire establishes the system's risk classification. In this questionnaire, it is determined whether an additional responsibility along the AI value chain applies.

Role and status **Q3 – Is the AI system already in use?**

Q3 determines the status of the AI system, which controls the compliance deadlines that apply to the user's role.

If the AI system is "In use", the user must comply with the requirements of the AI Act by 2030 at the latest (Article 111 AI Act). If the AI system is "In development", shorter, type-dependent deadlines apply (Article 113 AI Act, read in conjunction with Recitals 177-179).

After Q3 the questionnaire presents the result: the role(s) identified for the user, the status of the AI system. Responsibilities depend on the risk category of the AI system. It is therefore suggestion to users to fill the Risk category questionnaire.

Flowchart Role and status

Algorithm Audit

Flowchart – Role and status (Art. 2, 3, 25, 43, 111, 113 AI Act)

This schematic representation shows the logic required to determine what your role is in relation to the AI system and how the usage status of the system can be determined according to Article 2, 3, 25, 43, 111 and 113 of the AI Act. The complete questions can be found in the AI AQT tool itself.

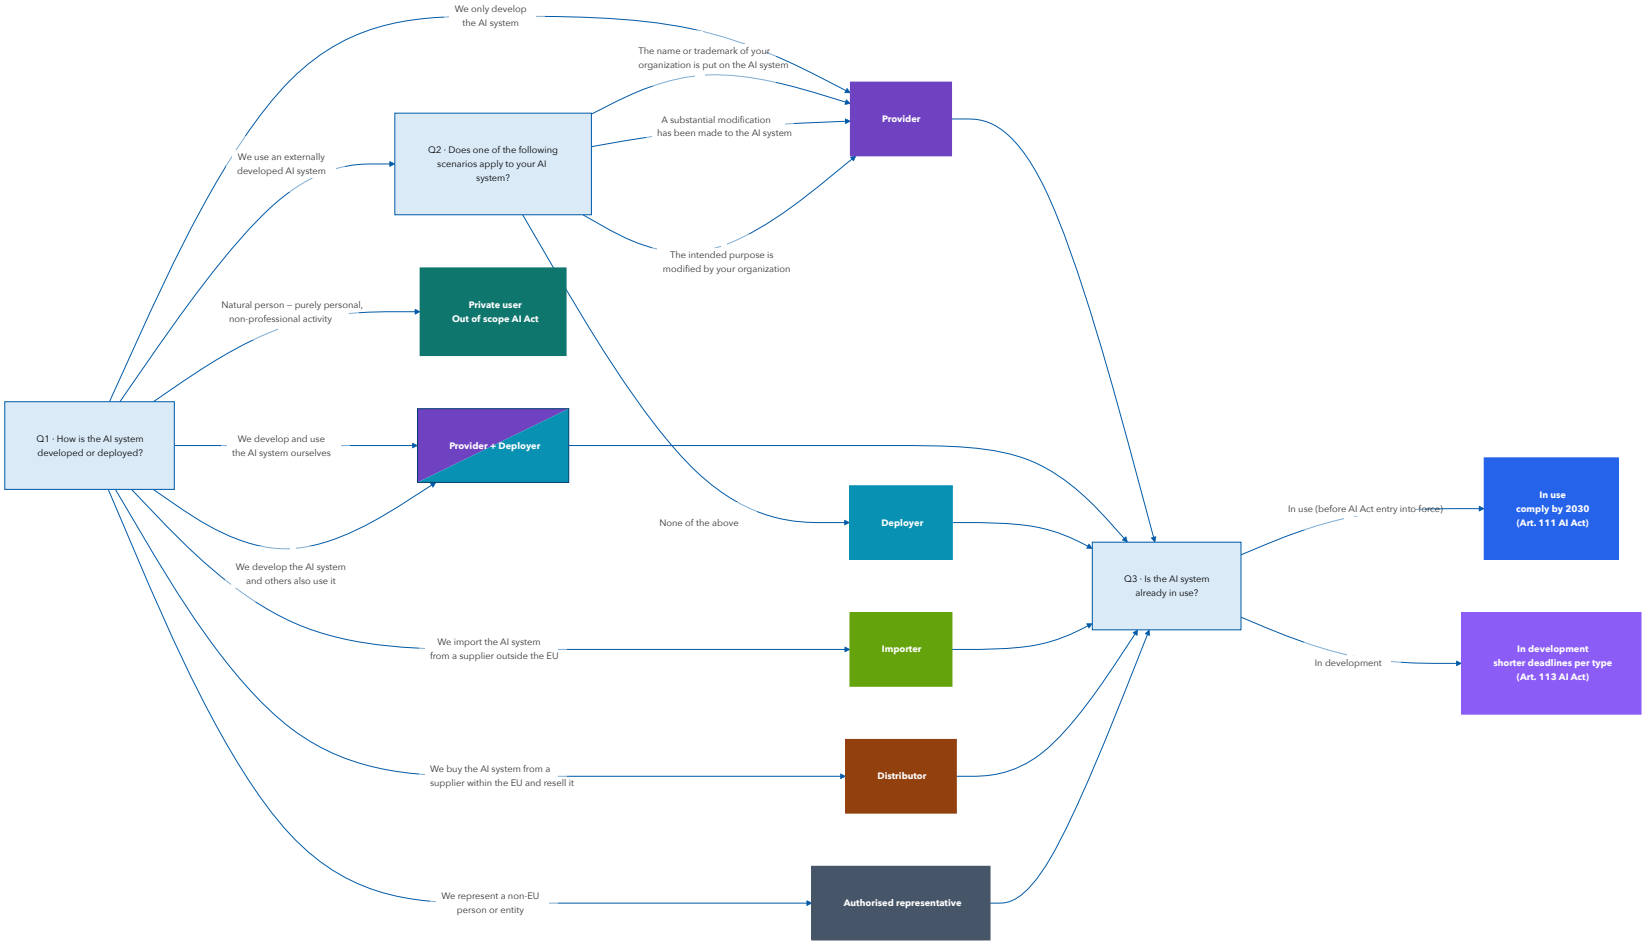

Figure 27 - Flowchart determining your role in relation to the AI system and status of the AI system.

6. Questionnaire Risk category

The AI Act of the European Union (EU) follows a risk-based approach: Depending on the risk category an AI system falls into different obligations apply (see section 7). In this section the risk category is determined. Questionnaire Risk category distinguishes five categories:

- **No requirements** – No obligations under the AI Act apply to the AI system (Articles 5 and 6). Obligations under other legislation (such as the GDPR) may still apply.
- **Prohibited AI systems** – The AI system meets the definition of a prohibited AI practice and may not be placed on the market, put into service or used within the Union (Article 5).
- **High-risk AI systems** – The AI system meets the classification rules for high-risk AI systems and must comply with the requirements depending on the role of your organisation in relation to the AI system. See section 7.
  Generative and interactive AI – Article 50 imposes specific transparency obligations on providers and deployers of generative or interactive AI systems. These obligations depend on the type of AI used. See section 7. These obligations apply regardless of the high-risk classification.
- **Exception** – The AI system would otherwise be high-risk or prohibited, but an exception for defence, military use or national security (Article 2(3)), scientific research and development (Article 2(6) and 2(8)) or narrow procedural/preparatory tasks (Article 6(3)) may apply. Verification with legal experts is required to confirm whether the exception is available.

Questionnaire Risk category determines, on the basis of 5-15 questions, into which of the above categories an AI system falls. The questionnaire opens with two gate questions that decide which branch is followed. First, it is checked whether Annex I – existing EU product-safety laws that are relevant for classifying AI systems as high-risk – is applicable (Q1). Then, it is checked whether the AI system uses biometric data (Q2). Both questions need to be answered before the questionnaire branches into more specific assessments.

Section 6.1 explains how generative and interactive AI are identified. The remaining sections describe the questions in the order in which they are asked: questions about Annex I (6.2), the biometrics branch (6.3), transparency obligations and non-consensual fake nude imagery (6.4), nine high-risk domains from Annex III (6.5), horizontal Article 5 prohibitions (6.6), Article 6(3) profiling and limited-task exceptions (6.7) and Article 2 exceptions (6.8). This chapter concludes with the overall flowchart of the Risk category questionnaire.

6.1 Annex I – List of Union harmonisation legislation

Q1 – Is the AI system intended to be used as a safety component of a product that falls under Annex I of the AI Act?

Article 6(1) classifies an AI system as high-risk if it is intended to be used as a safety component of a product covered by EU harmonisation legislation listed in Annex I of the AI Act, or if the AI system is itself such a product, provided that the product must undergo a third-party conformity assessment under that legislation. Section A of Annex I lists the modern EU product-safety regimes built around the New Legislative Framework (NLF) – a standardised approach to CE marking, conformity assessment, accreditation, and market surveillance.[32]

Annex I – Section A:

1. Toy Safety Directive (Directive 2009/48/EC)
2. Recreational Craft Directive (Directive 2013/53/EU)
3. Lifts Directive (Directive 2014/33/EU)

[32] The 'Blue Guide' on the implementation of EU product rules, European Commission (2022).

4. ATEX Directive – equipment for explosive atmospheres (Directive 2014/34/EU)
5. Radio Equipment Directive (Directive 2014/53/EU)
6. Pressure Equipment Directive (Directive 2014/68/EU)
7. Cableway Installations Regulation (Directive 2016/424)
8. Personal Protective Equipment Regulation (Directive 2016/425)
9. Gas Appliances Regulation (Directive 2016/426)
10. Medical Devices Regulation (Directive 2017/745)
11. In Vitro Diagnostic Medical Devices Regulation (Directive 2017/746).

Systems that relate to a Section A product are subject to the full set of high-risk obligations (see section 7). These obligations are integrated into the existing sectoral conformity-assessment procedures, so manufacturers typically go through a single combined assessment rather than two parallel ones.

Section B lists sector-specific regimes that fall outside the NLF, mainly covering transport and aviation. The AI Act imposes no direct obligations for providers or deployers of these systems. Their duties come from the operator-side rules of the sectoral framework.

Annex I – Section B:

1. Civil aviation security (Regulation 300/2008)
2. Two- and three-wheel vehicles and quadricycles (Regulation 168/2013)
3. Agricultural and forestry vehicles (Regulation 167/2013)
4. Marine equipment (Directive 2014/90/EU)
5. Rail system interoperability (Directive (EU) 2016/797)
6. Motor vehicle type-approval (Regulation (EU) 2018/858 and Regulation (EU) 2019/2144)
7. Civil aviation / EASA rules, insofar as they concern unmanned aircraft (Regulation (EU) 2018/1139).

When the user selects 'None of the below', the questionnaire continues to the biometric data screen (Q3). When the user selects an Annex I directive or regulation, the questionnaire presents Q2, which asks whether the product requires a third-party conformity assessment under that legislation. See Figure 28.

Q2 – Is the product whose safety component is the AI system, or the AI system itself as a product, required to undergo a third-party conformity assessment under EU harmonisation legislation?

Under Article 6(1), an AI system is classified as high-risk through the "safety component" route only if the same Annex I legislation selected in Q1 also requires the product (or the AI system itself, where it is the product) to undergo a third-party conformity assessment. This is what Q2 asks. See also Recital 50 and 51.

If the answer to Q2 is 'Yes' in combination with Annex I Section A legislation in Q1, obligations for high-risk AI systems apply (see section 7).

If the answer is 'Yes' in combination with Annex I Section B legislation in Q1, the AI system is classified as high-risk but is not subject to high-risk obligations. For these systems, AI Act requirements apply through existing sectoral regimes (Articles 102-109). Providers carry the compliance responsibility through the sectoral framework that already governs their product.

When the answer is 'No', the user is forwarded to the biometric data screen (Q3), as the AI system can still be high-risk under Annex III. See Figure 28.

Q3 – Does the AI system use biometric data?

According to Article 6(2), AI systems referred to in Annex III are also considered to be high-risk. Annex III(1) lists AI systems that use biometric data as a high-risk category. Biometric data are defined

as personal data resulting from specific technical processing related to the physical, physiological or behavioural characteristics of a natural person, which allow or confirm the unique identification of that person (Article 3(34)). See also Recital 14.

If the AI system uses biometric data ('Yes'), the user is forwarded to the biometrics branch (Q4). If the answer is 'No', the user is forwarded to the Article 50 transparency scenarios (Q11).

To assist users, the following remark is provided:

The key distinction of biometric data is: does the data measure a unique, stable, biological or behavioral trait of the body itself – or does it merely describe or record something associated with a person?

Examples of biometric data:
> Fingerprints, iris or retina patterns
> Facial and hand geometry, voice print, vein patterns, ear shape
> DNA
> Someone's unique way of typing (keystrokes) or walking (as captured by sensors or video analysis).

Examples of non-biometric data:
> A photo of someone
> Typed text or written signature
> Name or ID number
> Medical diagnosis
> Location or GPS data
> Passwords
> IP address or device fingerprints (tied to devices, not bodies)
> Age or birth date, height and weight.

6.2 Biometrics branch

Q4 – What does the AI system do with biometric data?

Q4 splits the biometrics branch into five sub-routes. Each answer option triggers a more specific check that targets a potentially prohibited practice under Article 5:
> "Creates or expands a facial recognition database." → Q5 (Article 5(1)(e)).
> "Remote identification of individuals." → Q6, then Q7 (Article 5(1)(h) and Annex III(1)).
> "For categorizing (labeling) people." → Q8 (Article 5(1)(g)).
> "Recognizes or infers emotions or intentions." → Q9 (Article 5(1)(f)).
> "Other purpose." → directly to the Article 50 transparency scenarios (Q11).

See also Recital 14 and 54.

Q5 – Does this involve large-scale collection of facial images from the internet or security cameras (scraping)?

Q5 is shown when Q4 is answered with 'Creates or expands a facial recognition database'. Article 5(1)(e) prohibits AI systems that create or expand facial recognition databases through the untargeted scraping of facial images from the internet or CCTV footage. See also Recital 43. If the answer is 'Yes', the questionnaire concludes that a prohibited practice is at stake and routes the user via the Article 2 exception screen to the Prohibited outcome. If the answer is 'No', the user is forwarded to the Article 50 transparency scenarios (Q11).

Q6 – Is the AI system only used to confirm that someone is who they say they are (verification or authentication)?

Q6 is shown if answer option 'Remote identification of individuals' is selected in Q4. If the answer is 'Yes', the AI system does not fall within the Article 5(1)(h) prohibition and the user is forwarded to the Article 50 transparency scenarios (Q11). If the answer is 'No', the user is forwarded to Q7 to assess the law-enforcement context. See also Recital 17 and 54.

To assist users, the following remark is provided:

Verification and authentication mean confirming an already-known identity: the person is present and willing, and the system checks them against a reference that belongs to them. If the person is unknown, passiveor not being matched against their own pre-stored reference, the system is not verifying or authenticating – it is trying to find out who someone is (identification) or infer something about them (categorisation).

Examples of verification and authentication:

> Unlocking a phone by matching your face to the face stored on the device
> Fingerprint reader giving access to building by matching your fingerprint to your enrolled profile
> A camera checks if your face matches the photo in your passport.

Examples not involving verification or authentication, but identification or categorization:

> A camera at a stadium entrance scans every face in the crowd to find a wanted suspect (identification)
> An airport uses facial recognition to spot known terrorists among passengers without their knowledge (identification)
> A shop uses cameras to detect the emotions or stress levels of customers (categorisation)
> An HR system analyses job applicants' facial expressions during video interviews (categorisation)
> A system checks whether someone walking past a camera is male or female based on their face (biometric categorisation)
> A police database is searched with CCTV footage to identify who committed a crime (identification).

Q7 – Is the AI system used for public safety, crime prevention, investigation or prosecution or the execution of criminal justice?

Q7 is only shown when Q6 is answered with 'No'. It operationalises Article 5(1)(h), which prohibits real-time remote biometric identification in publicly accessible spaces for the purpose of law enforcement (subject to narrow exceptions). See also Recital 32 and 33.

If the answer is 'Yes' to Q7, the questionnaire routes the user to the Article 50 transparency scenarios (Q11). If the answer is 'No', the user is forwarded to the cyber security and data protection exception (Q10).

Q8 – Does the AI system estimate race, ethnicity or origin, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation?

Q8 is shown when Q4 is answered with 'For categorizing (labeling) people'. Article 5(1)(g) prohibits the placing on the market, putting into service or use of biometric categorisation systems that categorise natural persons individually based on their biometric data to deduce or infer race, ethnic origin, political opinions, trade-union membership, religious or philosophical beliefs, sex life or sexual orientation. These types of characteristics are referred to a sensitive personal characteristics according to Article 9 GDPR. The lawful labelling or filtering of biometric datasets by law enforcement is not covered by this prohibition.

If the answer is 'Yes', the questionnaire routes the user via the Article 2 exception screen to the Prohibited outcome. If the answer is 'No', the user is directed to the cyber security and data protection exception (Q10).

Q9 – Is the emotion-recognition AI system used in the workplace or in education?

Q9 is shown when Q4 is answered with 'Recognizes or infers emotions or intentions'. Article 5(1)(f) prohibits AI systems used to infer emotions of a natural person in the workplace or in education institutions, except where the AI system is intended to be put in place or to be put into service for medical or safety reasons.

If the answer is 'Yes', the questionnaire routes the user via the Article 2 exception screen to the Prohibited outcome. If the answer is 'No', the user is sent to the cyber security and data protection exception (Q10).

To assist users, the following remark is provided:

The work environment also includes recruitment and selection, and the home work environment. Education also includes admission to education and online education.

Q10 – Is the system used solely for the purpose of enabling cybersecurity and personal data protection measures?

Q10 is only shown within the biometrics branch when 'No' has been answered to the following prohibited practices:

- Q6: The AI system (processing biometrics data) is not only used to confirm that someone is who they say they are (verification or authentication).
- Q8: The AI system (processing biometrics data) does not estimate race, ethnicity or origin, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation.
- Q9: The AI system (processing biometrics data) is not used in the workplace or in education.

AI systems that are solely used for the purpose of enabling cybersecurity and personal data protection measures are not considered to be high-risk AI systems. See Recital 54. This type of AI systems refer to systems whose only job is to defend IT systems or safeguard personal data. For example, tools that detect malware, block intrusions, filter phishing, control access, anonymize data or spot data leaks.

If Q10 is answered with 'Yes', the Recital 54 exception applies and the user is forwarded to the Article 50 transparency Scenarios (Q11).

If Q10 is answered with 'No', the Recital 54 exception doesn't apply and users are forwarded via the Article 2 exception screen to the Prohibited outcome.

6.3 Transparency obligations and prohibited non-consensual fake nude imagery

Q11 – Is the AI system used for non-consensual fake nude imagery?

Q11 is displayed in all cases where the user has not already been redirected to the exception screen prior to the determination that the AI system might be prohibited. Q11 determines whether generative or interactive AI is at stake. For these type of AI systems, additional transparency obligations apply (see section 7). The user selects every applicable scenario from a multi-select list (e.g., direct interaction with natural persons, synthetic-content generation, emotion or biometric-categorisation use, content generation or none of the above). The four scenarios refer to Article 50(1)-50(4).

If a scenario related to (synthetic) image generation is selected, the user is redirected to Q12 on non-consensual fake nude imagery. In all other cases, the user is forwarded to the Annex III domain gate (Q13). Selected scenarios add a transparency-obligation

badge to the results page. Selecting answer option 'None of the above' adds no such badge.

Q12 – Is the AI system used for non-consensual fake nude imagery?

Q12 is only shown when answer option 'It generates synthetic audio, image, video or text content.' or 'It generates or manipulates image, audio, video or public-interest text content.' is selected in Q11. In Q12, it is determined whether image generation is used for removing clothing from images of people without consent, a process also referred to as 'nudification' – which is prohibited under the AI Act as a result of AI Omnibus amendments.

If answered with 'Yes', the user is directly redirected to terminal node 'Prohibited'. If 'No', the user continues to Q12.

To assist users, the following remark is provided:

As clarified in Recital 134 and Article 50(6), the same AI system can be subject to Article 50 transparency obligations and be classified as high-risk under Article 6.

6.4 Annex III high-risk domains

Q13 – Does the AI system fall within any of the following domains?

Q13 is the entry point for the high-risk classification under Annex III(2) to Annex III(8). The user selects one of the below nine high-risk domains or 'None of the above'. Each domain triggers one or more domain-specific questions (described in the below sections). If none of the domains apply, the user is forwarded to horizontal Article 5 prohibitions (Q29).

- Critical physical and digital infrastructure →Q14 and Q15.
  - Tooltip: E.g., gas, water, heating, electricity supply; the management or operation of roads, bridges or road traffic; and critical digital infrastructure (internet nodes, DNS services, top-level domain name registration, cloud computing services, data center services, public electronic communication networks, and electronic communication services).
- Education or vocational training → Q16.
  - Tooltip: E.g., placement and admission, evaluation and testing, but also control such as plagiarism detection and anti-cheating software.
- Employment or HR → Q17, Q18 and Q19.
  - Tooltip: Examples how AI systems are used in the context of employment are:
    - Recording hours, measuring KPIs and performance
    - Task distribution
    - Recruitment and selection
    - Inflow and progression (employment conditions, contractual terms)
    - Outflow (termination of contract)
  - Also choose 'yes' when the application broadly falls under the above, even if the application plays a small role in an otherwise manual process.
- Government benefits or public services → Q20.
  - Tooltip: E.g., processes around benefits or allowances, such as application, administration, execution, enforcement, recovery, and objection. This also includes waiving fines, repayments, or granting a payment arrangement. Select this option if the application plays a small role in the process.
- Assessment of creditworthiness, life or health insurance, or emergency services → Q21.
  - Tooltip: For example, to determine if someone can get a payment arrangement for paying taxes, a fine, or repaying a benefit.
- Law enforcement → Q22 and Q23.
  - Tooltip: Law enforcement includes the prevention, investigation, and prosecution of criminal offences, the execution of sentences, and the protection of public safety. Think of police, special enforcement officers, and

the execution of (community) sentences by probation services. Other agencies, such as juvenile justice organizations and specific (inspection) services, can also engage in law enforcement. Also select this answer for organizations and institutions that support law enforcement agencies, such as a forensics institute, and for activities carried out on behalf of law enforcement agencies.

- Migration, asylum or border control → Q24, Q25 and Q26.
- Administration of justice or dispute resolution → Q27.
  - Tooltip: Professional matters, in which a judge makes a ruling, fall under the judiciary. Alternative dispute resolution includes disciplinary law, ombudsmen, arbitration, and complaints committees such as the Advertising Code Committee, the Council for Journalism, and the National Complaints Committee for Education.
- Elections or democratic processes → Q28.
- None of the above → horizontal Article 5 prohibitions (Q29).

Q14 – Does the AI system fulfill a safety function within the management or operation of this infrastructure?

Q14 is only shown when answer option 'Critical physical and digital infrastructure' is selected in Q13. Q14 asks whether the AI system fulfils a safety function within the management or operation of critical infrastructure such as gas, water, heating or electricity supply, road traffic, or critical digital infrastructure, as mentioned in Annex III(2). See also Recital 55.

If the answer is 'No', the AI system is not classified as high-risk under Annex III(2) and the user proceeds to the horizontal Article 5 prohibitions (Q29). If the answer is 'Yes', the user is redirected to Q15.

To assist users, the following remark is provided:

A component with a safety function is intended to eliminate or reduce a risk. If the failure or malfunction of the AI system endangers the health and safety of people or property, choose yes. Examples of components with a safety function are:

- A system that automatically closes a lane for a tunnel during road congestion to prevent traffic jams in the tunnel
- Fire alarm systems
- Other detection or alarm functions that respond to over/under pressure, temperatureor voltage
- Collision avoidance systems in vehicles that automatically brake to prevent a collision
- Firewall or antivirus software.

Q15 – Is the AI system a cybersecurity system?

Q15 is only shown when Q14 is answered with 'Yes'. Q15 asks whether the AI system is a cybersecurity system. Pursuant to Recital 55, cybersecurity systems supporting critical infrastructure are not classified as high-risk under Annex III(2). See also Recital 55. Therefore, a 'Yes' here exits the high-risk path and a 'No' forwards the user to Article 2 high-risk exceptions.

To assist users, the following remark is provided:

Also choose an option when the AI system plays a part in the process, without being the only factor.

Q16 – In which of the following education contexts is the AI system used?

Q16 is only shown when answer option 'Education or vocational training' is selected in Q13. Q16 lists four high-risk uses of AI in education and vocational training, as mentioned in Annex III(3): allocating or admitting students, assessing the level or quality of education someone receives, evaluating learning outcomes or guiding personalised learning, and monitoring or detecting unauthorised behaviour of

students (e.g., proctoring or anti-cheating software). See also Recital 56. Selecting any of these forwards the user to the Article 2 high-risk exceptions. Selecting 'None of the above' continues to the horizontal Article 5 prohibitions (Q29).

'Assessing the level or quality of education that someone receives or has access to' should also be selected if the application contributes to the assessment. There does not need to be (fully) automated assessment.

Examples of evaluating learning outcomes or guiding the learning process (personalized learning) are:

> Evaluation of (multiple choice) tests
> Grading assistance
> Assessment of learning pace or style
> Personalization of the curriculum or learning process.

To assist users, the following remark is provided:

Also choose an option when the AI system plays a part in the process. There does not need to be fully automated assessment or decisions.

Q17 – In which of the following employment areas is the AI system used?

Q17 is only shown when answer option 'Employment or HR' is selected in Q13. Q17 lists six high-risk uses in employment, as mentioned in Annex III(4): targeted job advertisements, analysing or filtering applications, assessing candidates, decisions on contractual terms or promotions, decisions on outflow or termination, and task allocation among employees. See also Recital 57. Selecting any of these forwards the user to the Article 2 high-risk exceptions. Selecting 'None of the above' continues to Q18.

Placing targeted job advertisements includes both targeted recruitment and writing targeted job descriptions. 'Employees' have a broad definition, including regular employees but also freelancers and employees hired through a third party (temporary, secondment). Determining the distribution of tasks among employees includes administrative processes where hours, productivity, speed, breaks, performance, or behavior are processed.

To assist users, the following remark is provided:

Also choose an option when the AI system plays a part in the process. There does not need to be fully automated assessment or decisions.

Q18 – Is the AI system used in processes where activities, behavior or performance of employees are processed?

Q18 is only shown when answer option 'None of the above' is selected in Q17. In contrast to the closed list of Q17, Q18 is broader formulated to capture tools, especially record-keeping HR systems, that don't fit the Q17 verbs (e.g. workforce-analytics, productivity-tracking, shift/scheduling tools, sentiment dashboards). Recital 57 follows the same split. If 'Yes', the AI system might be high-risk and users are brought to Q20 to verify this. If 'No', users are redirected to horizontal Article 5 prohibitions (Q29).

To assist users, the following remark is provided:

> Think of administrative processes where hours, productivity, speed, breaks, performanceor behavior are processed.
> We mean employees in the broadest sense, including employees but also freelancers and employees hired through a third party (temporary, secondment).

Q19 – Is there evaluation or monitoring of behavior or performance in this process or in related processes?

Q19 is only shown when question Q18 is answered with 'Yes'. Annex III(4)(b) requires both Q18 and Q19 to be true: processing employees' activities and monitoring or evaluating them. Splitting Q18 from Q19 prevents pure record-keeping HR systems from being mis-classified as high-risk. See also Recital 57. If 'Yes', the AI system might be high risk and users are brought to the Article 2 high-risk exceptions. If 'No', users are redirected to horizontal Article 5 prohibitions (Q29).

To assist users, the following remark is provided:

Think of monitoring productivity or speed or other performance.

Q20 – Is the AI system used to assess whether someone is eligible for a service, provision or benefit or to decide whether a service is granted, denied, limited, withdrawn or reclaimed?

Q20 is only shown when answer option 'Government benefits or public services' is selected in Q13. Q20 asks whether the AI system is used to assess eligibility for a public service, provision or benefit, or to decide whether such a service is granted, denied, limited, withdrawn or reclaimed, as listed in Annex III(5)(a). See also Recital 58. A 'Yes' forwards the user to the Article 2 high-risk exceptions. A 'No' continues to the horizontal Article 5 prohibitions (Q29).

To assist users, the following remark is provided:

Emergency services include police, fire brigade, and ambulance. Credit scoring also includes determining if someone can get a payment arrangement for taxes, fines or repaying a benefit.

Q21 – Is the AI system used for any of the following?

Q21 is only shown when answer option 'Assessment of creditworthiness, life and health insurances or emergency services' is selected in Q13. Q21 lists three high-risk uses: creditworthiness assessment or credit scoring of individuals, risk assessment or pricing for life or health insurance, and evaluating, deploying or prioritising emergency services or triage of urgent medical care, as listed in Annex III(5)(b)-(d) and Recital 58. Selecting any of these forwards the user to the Article 2 high-risk exceptions. Selecting 'None of the above' continues to the horizontal Article 5 prohibitions (Q29).

To assist users, the following remark is provided:

Emergency services include police, fire brigade, and ambulance. Credit scoring also includes determining if someone can get a payment arrangement for taxes, fines or repaying a benefit.

Q22 – Is the AI system used in assessing the risk that someone will (re)commit a criminal offense?

Q22 is only shown when answer option 'Law enforcement' is selected in Q13. Q22 asks whether the AI system is used to assess the risk that someone will (re)commit a criminal offence. See also Recital 42 and Recital 59. A 'Yes' routes the user to horizontal Article 5 prohibitions (Q29). A 'No' continues to Q23.

To assist users, the following remark is provided:

Also think of (administrative) offenses and fraud risk scores, where an estimate is made of whether someone might commit a criminal offense and an estimate of the likelihood of recidivism.

Q23 – Is the AI system used in assessing the risk that someone will (re)commit a criminal offense?

Q23 is only shown when answer Q22 is answered with 'No'. Q23 lists five high-risk uses in law enforcement, as stated in Annex III(6): assessing or estimating that someone will become a victim of crime; acting as a polygraph or similar instrument; assessing the reliability of evidence; automatically evaluating or predicting characteristics, behaviour or personal aspects in the context of crime; and assessing personality traits, characteristics or previous criminal behaviour. See also Recital 59. Any of these forwards the user to horizontal Article 5 prohibitions (Q29).

To assist users, the following remark is provided:

Also choose an option when the AI system plays a part in the process.

Q24 – Does the AI system do any of the following?

Q24 is only shown when answer option 'Migration, asylum or border control' is selected in Q13. Q24 lists three high-risk uses, as stated in Annex III(7): acting as a polygraph or similar instrument; assessing risks for a person entering EU territory; and processing asylum, visa or residence-permit applications or related complaints. See also Recital 60. Any of these forwards the user to the Article 2 high-risk exceptions. 'None of the above' continues to Q25.

Processing asylum, visa or residence permit applications or handle complaints related to the eligibility of the applicant is also applicable when the AI system is used to detect, recognise or identify individuals, or for assessing the reliability of evidence for applications or complaints.

Q25 – Is the AI system used for the detection, recognition or identification of individuals?

Q25 is only shown when answer option 'None of the above' is selected in Q24. Q25 asks whether the AI system is used for detection, recognition or identification of individuals in a migration or border-control context. A 'No' exits the high-risk path and the user is redirected to horizontal Article 5 prohibitions (Q29). See Annex III(7) and Recital 60. If 'Yes', whether a high-risk classification of the AI systems is applicable is determined in Q26.

Q26 – Is the AI system solely used to verify a passport or other travel document to confirm that a specific person is who they claim to be?

Q26 is only shown when Q25 is answered with 'Yes'. Q26 asks whether the AI system is used solely to verify a passport or other travel document. A 'Yes' here brings users to horizontal Article 5 prohibitions (Q29). A 'No' forwards the user to the Article 2 high-risk exceptions. See Annex III(7) and Recital 60.

To assist users, the following remark is provided:

Choose 'no' when the AI system is also used for identification that is broader than just checking a passport or other travel document.

Q27 – Is the AI system used only for supporting administrative activities that do not affect the judiciary in individual cases?

Q27 is only shown when answer option 'Administration of justice or dispute resolution' is selected in Q13. Q27 asks whether the AI system is used only for supporting administrative activities that do not affect the judiciary in individual cases (such as anonymising decisions or supporting communication between staff), as stated in Annex III(8)(a). See also Recital 61. A 'Yes' brings users to horizontal Article 5 prohibitions (Q29). A 'No' forwards the user to the Article 2 high-risk exceptions.

To assist users, the following remark is provided:

Think of anonymizing or pseudonymizing judicial decisions, documents or data, communication between staff members, and administrative tasks.

Q28 – Is the AI system used within processes for elections or referenda in a way that can influence the outcome or voting behavior of people?

Q28 is only shown when answer option 'Elections or democratic processes' is selected in Q13. Q28 asks whether the AI system is used in elections or referenda in a way that can influence the outcome or the voting behaviour of natural persons, as stated in Annex III(8)(b). A 'Yes' forwards the user to the Article 2 high-risk exceptions. A 'No' brings users to horizontal Article 5 prohibitions (Q29). See also Recital 62.

To assist users, the following remark is provided:

This does not include AI systems to the output of which natural persons are not directly exposed, such as tools used to organise, optimise or structure political campaigns from an administrative or logistical point of view.

6.5 Article 5 horizontal prohibitions

Q29 – Does the AI system use techniques people cannot consciously perceive (e.g., subliminal techniques), or is it designed — or known to unintentionally — manipulate people's behavior or choices without their awareness?

Q29 is shown if an AI system has exited the high-risk path but could still potentially qualify as a prohibited application under Article 5. Under Article 5(1)(a), an AI system is prohibited if it uses techniques people cannot consciously perceive – such as briefly displayed subliminal images – or is designed (or known) to manipulate people's behaviour or choices without their awareness, whether intentionally or not.

If the answer to Q29 is 'No', the user is forwarded to Q31. If the answer is 'Yes', Q30 asks whether the AI system can potentially have negative consequences for the user or others, which is a further requirement under Article 5(1)(a) for the system to qualify as prohibited.

To assist users, the following remark is provided:

Subliminal techniques include showing images so briefly that they cannot be consciously perceived. Manipulation without awareness includes AI chatbots that give advice influencing decisions, applications that stimulate addictionor systems that prevent users from filing a complaint while they intended to.

Q30 – Can the AI system potentially have negative consequences for the user or other people?

Q30 is shown if an AI system Q29 is answered 'Yes'. Q30 is the second-stage test whether the potential prohibition of the AI system under Article 5(1)(a) is applicable, i.e., distort behaviour in a manner that "causes or is reasonably likely to cause significant harm". A 'Yes' to Q30 routes the user to Article 2 exceptions. A 'No' brings the user to Q31.

To assist users, the following remark is provided:

Think of an AI health chatbot intended to give tips on a healthy lifestyle and tailored advice for psychological and physical exercises. However, if the chatbot encourages unhealthy habits or dangerous activities (e.g., excessive exercise without rest or drinking water) that they would not have done otherwise, it has negative consequences for the user.

Q31 – Does the AI system evaluate, score or classify individuals based on their social behavior or personal characteristics in a way that could lead to unfavorable treatment?

Q31 is shown when Q29 or Q30 are answered with 'No', i.e., no prohibited subliminal or manipulative techniques are applicable. Q31 asks whether the AI system might stille be prohibited under Article 5(1)(c) – also referred to as 'social scoring'. Whether social scoring is applicable is also a two-step exercise. If Q31 is answered with 'No', the questionnaire concludes that 'No requirements' follow from the AI Act. If Q31 is answered with 'Yes', users are directed to a follow up question (Q32) to determine whether prohibited social scoring is applicable.

To assist users, the following remark is provided:

Examples of social behavior are:

> Driving behavior
> Communication style
> Breaking (un)written rules

Examples of personal characteristics are:

> Age
> Height
> Gender
> Economic situation
> Movement patterns
> Interests

Q32 – Does the unfavorable treatment occur in a completely different context from the behavior or characteristic it is based on or is it disproportionate to that behavior?

Q32 is shown if Q31 is answered with 'Yes'. Q32 checks whether the AI system leads to detrimental or unfavourable treatment "in a context unrelated to the data origin" or "unjustified or disproportionate to the gravity", which is a condition of the prohibition of Article 5(1)(c). A 'Yes' to Q32 routes users to Article 2 exceptions. A 'No' concludes that 'No requirements' follow from the AI Act.

To assist users, the following remark is provided:

> Unfavorable treatment in an unrelated context is, for example, a fraud risk score that leads to an investigation, where the fraud risk score is based on characteristics without clear relevance to the assessment of fraud, such as having a partner with a certain nationality or behavior on social platforms.
> Disproportionate treatment is, for example, someone getting a higher fraud risk score because municipal taxes were once paid late or a student not being admitted to further education because they were once caught cheating.]

6.6 Article 6(3) – Profiling and limited tasks

Q33 – Is the AI system used to automatically evaluate or predict characteristics, behavior or personal aspects?

Q33 is only shown when AI systems are potentially high-risk and no exceptions apply ('None of the above' is selected for the Article 2 exception question). Q33 asks whether the AI system is used to automatically evaluate or predict characteristics, behaviour or personal aspects of natural persons (i.e., profiling within the meaning of Article 4(4) GDPR). Article 6(3) provides that an AI system referred to in Annex III is not high-risk if it does not pose a significant risk of harm to the health, safety or fundamental rights of natural persons. However, this carve-out is not available when the AI system performs profiling. See also Recital 53. A 'Yes' to Q33 therefore yields the high-risk outcome. A 'No', users are redirected to Q34.

To assist users, the following remark is provided:

Think of automatically evaluating professional performance, economic situation, health, personal

preferences, interests, reliability, behavior, locationor movements of a person to analyze or predict.

Q34 – Is the AI system only used for narrow procedural tasks or preparatory tasks without influencing the outcomes of a (decision-making) process?

Q34 is only shown when Q33 is answered with 'No' (profiling is not applicable). Q34 asks whether the AI system is used only for narrow procedural or preparatory tasks: a narrow procedural task; the improvement of the result of a previously human-performed activity; the detection of decision-making patterns or deviations from previous decision-making patterns; or a preparatory step that does not influence the outcome. As specified in Article 6(3) and Recital 53. Selecting any of these task categories yields the high-risk-with-exception outcome (Article 6(3)). Selecting 'No' yields the high-risk outcome.

6.7 Article 2 exception

Article 2 exception – Does one of the following exceptions apply to your application?

The Article 2 exception question asks whether the AI system falls under one of the Article 2 exceptions: scientific research and development (Article 2(6)), defence or military purposes (Article 2(3)), or national-security purposes (Article 2(3)). When the AI system is used for operational processes or regular business operations, the user should select 'None of the above'.

If a research exception is selected, the user is asked to describe the research (purpose, department, partners). The outcome in all three exception cases is the corresponding high-risk-with-exception or prohibited-with-exception, depending on the path that led to the Article 2 exception question. If 'None of the above' is selected and the user is on the high-risk path, the questionnaire continues to Q33 (Article 6(3) profiling check). If the user is on the prohibited path, 'None of the above' yields the Prohibited outcome.

Flowchart Risk classification

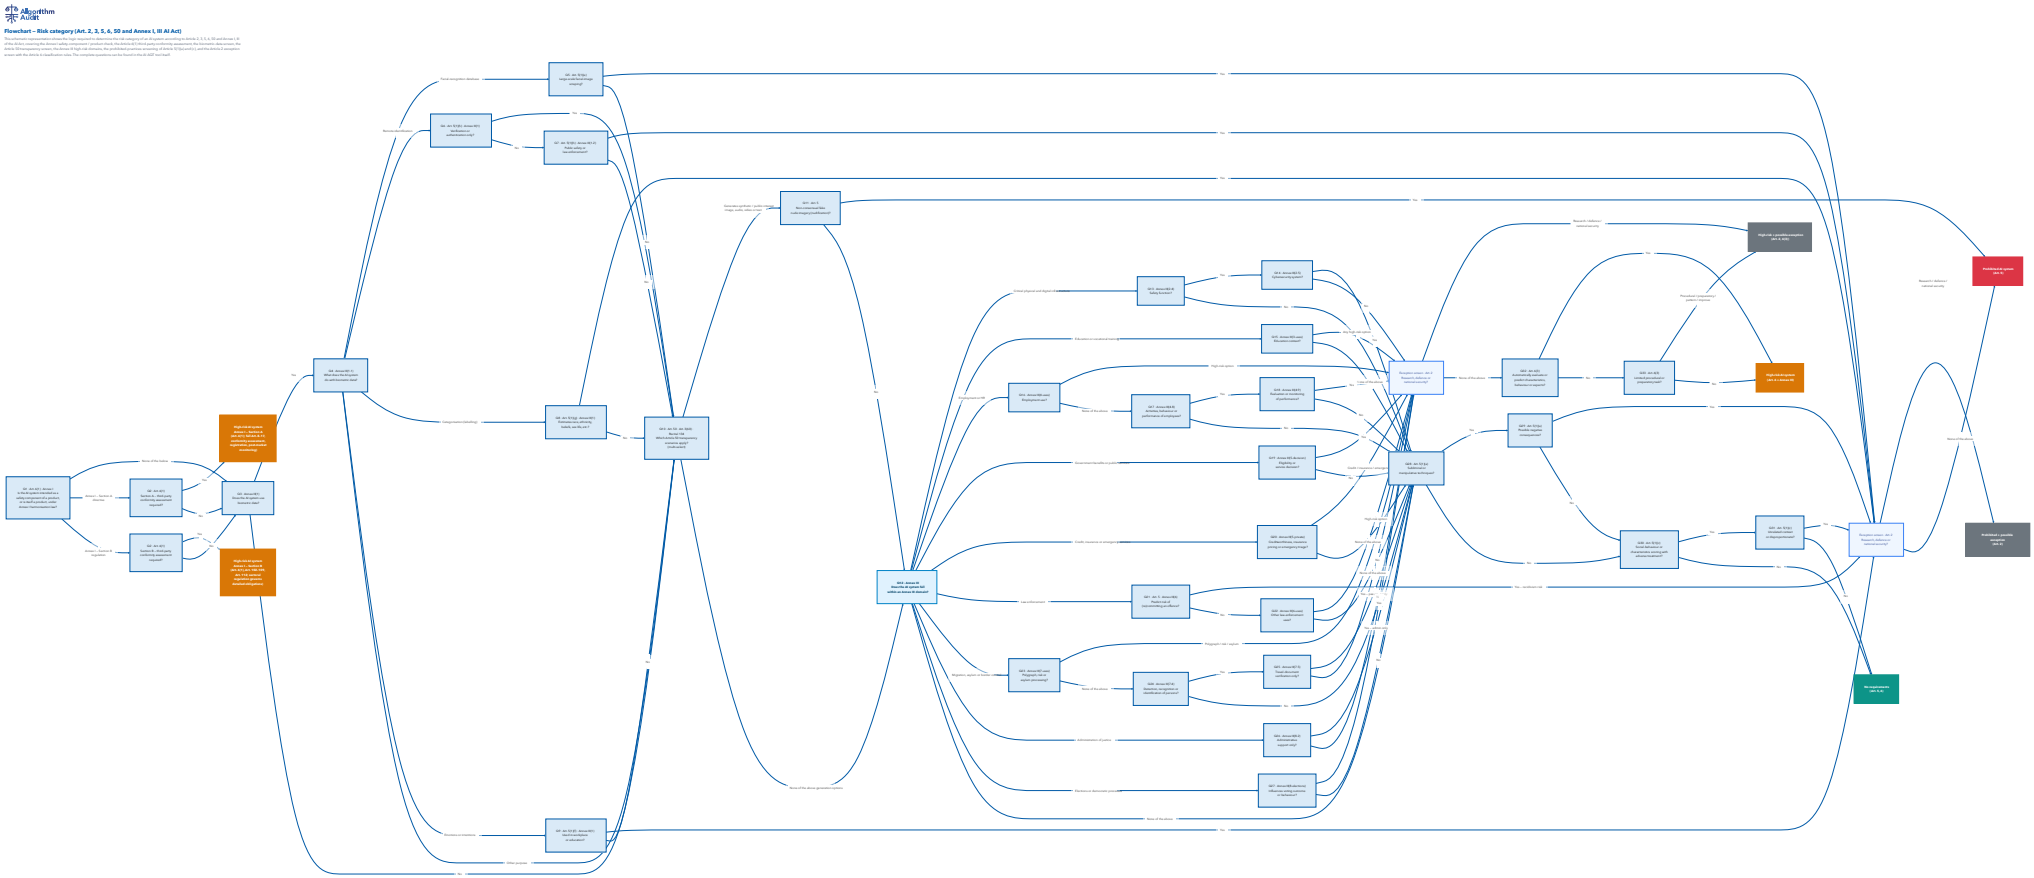

Figure 28 - Flowchart determining the risk category of the AI system.

7.Questionnaire Obligations

The questionnaire Obligations describes the requirements that follow from the AI Act depending on the role, status and risk category of the AI system. The questionnaire consists of three questions. When the questionnaires Role and status and Risk category have been completed, the answers are pre-filled automatically; they can also be selected manually.

- **Q1 – What role do you have in relation to the AI system?** The role determines which set of obligations applies. The AI Act distinguishes provider, deployer, importer, distributor and authorised representative. A private user, using an AI system in the course of a purely personal, non-professional activity, has no obligations under the AI Act. The roles provider and deployer can be held simultaneously; the other roles are treated as mutually exclusive.
- **Q2 – What is the usage status of the AI system?** The status (in use or in development) determines the compliance deadline that applies.
- **Q3 – What is the risk category of the AI system?** The risk category determines the substantive obligations. The categories are no requirements, prohibited, high-risk, generative and interactive AI, and exception. The categories high-risk and generative and interactive AI can apply at the same time.

The obligations follow from Chapter III of the AI Act (Article 6, read in combination with Annex I or Annex III) and from Article 50 of the AI Act. For AI systems classified as posing no requirements, no obligations follow from the AI Act. Where an exception applies, for example for scientific research and development as referred to in Article 2, the obligations are those of the underlying category to which the exception attaches. The sections below describe the obligations per risk category: prohibited AI systems (section 7.1), high-risk AI systems (section 7.2) and generative and interactive AI (section 7.3).

7.1 Obligations for prohibited AI systems

An AI system that meets the definition of a prohibited AI practice (Article 5) may not be placed on the market, put into service or used within the European Union. This prohibition has applied since 2 February 2025. There is no compliance pathway for prohibited AI systems: the system must be discontinued.

Provider

This AI system must be taken out of service immediately. Its intended use constitutes a prohibited AI practice under the EU AI Act and may not be placed on the market, put into service or used within the European Union.

Deployer

Stop using the AI system. Inform the provider that the system must be taken out of service immediately. Its intended use constitutes a prohibited AI practice under the EU AI Act and may not be placed on the market, put into service or used within the European Union.

Further guidance is available in the Guidelines on Prohibited Artificial Intelligence[33] published by the European Commission.

7.2 Obligations for high-risk AI systems

High-risk AI systems must comply with the requirements set out in Chapter III of the AI Act. Standalone high-risk AI systems – high-risk because they fall within an Annex III domain – must comply by 2 December 2027.[34] High-risk AI systems that are,

[33] Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689.
[34] Digital Omnibus on AI, European Commission (2026).

or are a safety component of, a product covered by Annex I must comply by 2 August 2028. For high-risk AI systems already in use, the transitional deadlines of Article 111 AI Act apply: systems placed on the market or put into service before 2 August 2026 that are intended to be used by public authorities must comply by 2 August 2030, while other standalone high-risk systems must comply by 2 August 2027.

The obligations differ per role. The most extensive set of obligations applies to the provider, the deployer, importer and distributor. Each carrying a more limited set. The relevant legal bases are Article 16 (provider), Articles 26 and 27 (deployer), Article 23 (importer) and Article 24 (distributor) of the AI Act.

A special case applies where the AI system is high-risk because it falls under Annex I Section B sectoral legislation – covering, among others, civil aviation security, motor vehicles and rail interoperability – that requires a third-party conformity assessment. In that case the AI Act does not impose new obligations beyond the sectoral framework. Obligations of Articles 8-17 are integrated into the relevant sectoral regulation (see Article 2(2)).

Provider

The provider of a high-risk AI system carries the most extensive obligations. Pursuant to Article 16, the provider must satisfy the following requirements.

- **Risk management system –** Establish, implement, document and maintain a risk management system throughout the entire lifecycle of the high-risk AI system (Article 9).
- **Data and data governance –** Ensure that training, validation and testing data sets are relevant, sufficiently representative and, to the best extent possible, free of errors (Article 10).
- **Record-keeping and logging –** Retain automatically generated logs for at least six months, or as otherwise required by law (Article 12).
- **Transparency and provision of information to deployers –** Provide clear instructions for use to deployers (Article 13).
- **Human oversight –** Build in technical measures enabling effective human oversight and meaningful human intervention (Article 14).
- **Accuracy, robustness and cybersecurity –** Design and develop the system to achieve an appropriate level of accuracy, robustness and cybersecurity throughout its lifecycle (Article 15).
- **Quality management system –** Establish and maintain a quality management system, including post-market monitoring (Article 17).
- **Technical documentation –** Prepare and maintain full technical documentation (Annex IV) before placing the system on the market (Article 11).
- **Conformity assessment –** Carry out the applicable conformity assessment procedure – self-assessment or third-party notified body, depending on the system type (Article 43).
- **CE marking** – Affix the CE marking to indicate conformity with the AI Act (Article 48).
- **EU Declaration of Conformity –** Draw up and sign the declaration (Article 47).
- **Registration –** Register the system in the EU database before placing it on the market (Article 49).
- **AI literacy –** Ensure that staff have an adequate level of AI literacy (Article 4).

Deployer

Pursuant to Articles 26 and 27, the deployer of a high-risk AI system must satisfy the following requirements.

- **Risk management system –** Monitor the operation of the system on the basis of the instructions for use and report identified risks or serious incidents to the provider, distributor and market surveillance authority without undue delay (Article 26(5)).

- **Data and data governance –** Where the deployer controls the input data, ensure it is relevant and sufficiently representative in view of the intended purpose (Article 26(4)).
- **Record-keeping and logging –** Keep the automatically generated logs for at least six months (Article 26(6)).
- **Transparency and instructions for use –** Use the AI system in accordance with the instructions for use and take appropriate technical and organisational measures to ensure compliance (Article 26(1)).
- **Human oversight –** Assign human oversight to persons with the necessary competence, training and authority (Article 26(2)).
- **Workers –** Inform workers' representatives and affected workers before putting the system into service at the workplace (Article 26(7)).
- **Transparency to affected persons –** Inform natural persons that they are subject to a high-risk AI system (Article 26(11)).
- **Fundamental rights impact assessment –** Carry out a fundamental rights impact assessment before deployment; this applies to public bodies and to certain private entities deploying systems listed in Annex III (Article 27).

Importer

Pursuant to Article 23, the importer of a high-risk AI system must:

- Verify, before placing the system on the market, that the provider has carried out the conformity assessment, drawn up the technical documentation and affixed the CE marking.
- Verify that the provider has appointed an authorised representative, where required.
- Ensure the AI system bears the required CE marking and is accompanied by the EU Declaration of Conformity and instructions for use.
- Refrain from placing a non-compliant system on the market.
- Maintain a copy of the EU Declaration of Conformity for ten years.
- Report any non-compliance or serious risk to the provider and competent authorities.
- Cooperate with national authorities and provide information upon request.

Distributor

Pursuant to Article 24, the distributor of a high-risk AI system must:

- Verify, before distributing the system, that it bears the required CE marking, is accompanied by the EU Declaration of Conformity and instructions for use, and that the provider and importer obligations have been met.
- Refrain from distributing a non-compliant system.
- Ensure that storage and transport conditions do not affect the system's conformity.
- Report non-compliance to the provider or importer, and, where a risk is present, to the competent authorities.
- Take corrective action, such as recall or withdrawal, where necessary after distribution.
- Cooperate with national authorities upon request.

7.3 Obligations for generative and interactive AI

Article 50 imposes transparency obligations on providers and deployers of certain AI systems, irrespective of whether those systems are also classified as high-risk. The deadline for the Article 50 transparency requirements is 2 August 2026. Only the deadline for the marking (watermarking) requirements under Article 50(2) is extended to 2 December 2026.

Article 50 distinguishes four sub-cases:

- **Article 50(1) –** AI systems interacting directly with natural persons.
- **Article 50(2) –** AI systems generating synthetic audio, image, video or text content.

- **Article 50(3)** – Emotion recognition and biometric categorisation systems.
- **Article 50(4) –** AI systems generating or manipulating deep fakes and AI-generated text published to inform the public on matters of public interest.

The transparency obligations differ per role and per sub-case. The information must always be provided in a clear and distinguishable manner, at the latest at the time of the first interaction or exposure (Article 50(5)). The sub-sections below set out the obligations for providers and deployers.

Provider

Interacting directly with natural persons (Article 50(1))

The provider must:

- Design and develop the AI system so that natural persons are clearly informed that they are interacting with an AI system, unless this is obvious from the perspective of a reasonably well-informed natural person, taking into account the circumstances and context of use.
- Embed the notification in the design of the system and provide it at the latest at the time of the first interaction (Article 50(5));
- Ensure the information is clear, distinguishable, accessible – including for vulnerable groups such as children, the elderly and persons with disabilities – and given in a context-appropriate manner.

Generating synthetic content (Article 50(2))

The provider must:

- **Marking –** Mark the outputs in a machine-readable format, for example through watermarks, metadata, cryptographic provenance methods or fingerprints.
- **Detection –** Make means available to detect that the output is AI-generated or manipulated, with human-readable results.

Both the marking and detection elements must be present. Marking alone is insufficient. The technical solution must be effective, reliable, robust and interoperable.

Emotion recognition and biometric categorisation (Article 50(3))

The provider must inform exposed persons that the system is in use and process personal data in accordance with the GDPR.

Deep fakes and AI-generated public-interest text (Article 50(4))

Apply to deployers and are set out below.

Deployer

Interacting directly with natural persons (Article 50(1))

The deployer must:

- Verify that the provider has implemented the disclosure of interactive AI, for example an opening line from a chatbot, a clear "AI assistant" label, or an equivalent signal.
- Ensure that its own configuration, branding or integration of the system does not undermine or remove the provider's disclosure. Where the deployer fine-tunes, rebrands or substantially modifies the system, Article 25 may convert the deployer into a provider, importing the full Article 50(1) provider obligations.

Deep fakes and public-interest text (Articles 50(2) and 50(4))

The deployer must:

- Disclose that content constituting a deep fake (as defined in Article 3(60)) has been artificially generated or manipulated.
- Disclose that AI-generated or manipulated text published with the purpose of informing the public on matters of public interest has been artificially generated or manipulated.

Timing: This information must be provided to users at the latest at the time of the first exposure to the system.

Exceptions: The disclosure obligation does not apply where the use is authorised by law to detect, prevent, investigate or prosecute criminal offences; where the content forms part of an evidently artistic, creative, satirical, fictional or analogous work or programme – in which case the obligation is limited to disclosing the existence of generated or manipulated content in a manner that does not hamper the display or enjoyment of the work; or where the AI-generated content has undergone a process of human review or editorial control and a natural or legal person holds editorial responsibility for the publication.

Emotion recognition and biometric categorisation (Article 50(3))

The deployer must:

- Inform all natural persons exposed to the system that it is being operated or used.
- Provide the information in a clear, distinguishable and accessible manner (Article 50(5)).
- The information may be given in writing, via standardised icons, orally, electronically, or in any combination, depending on the deployment context.

Timing: Provide the information to users at the latest at the time of the first exposure to the system.

Exception: systems permitted by law to detect, prevent or investigate criminal offences, with appropriate safeguards under Union law, are not subject to this obligation.

For Article 50(1) and Article 50(2), the European Commission's draft Guidelines on the implementation of the transparency obligations under Article 50 provide further detail on acceptable and insufficient disclosure techniques and on the quality requirements for marking and detection solutions.

Flowchart Obligations

Flowchart – Obligations by risk category, role and Article 50 scenario (Art. 4, 5, 9, 10-17, 26, 27, 43, 47- 50, 72 and Annex III, IV AI Act)

This schematic representation shows the obligations that apply, depending on the role of the actor (provider or deployer) and the risk category (prohibited, high-risk or generative and interactive AI), mirroring the conditional logic of the "Risk category results" page based on the Risk category questionnaire and Role and status questionnaire. The complete questions and result page text can be found in the AI AQT itself.

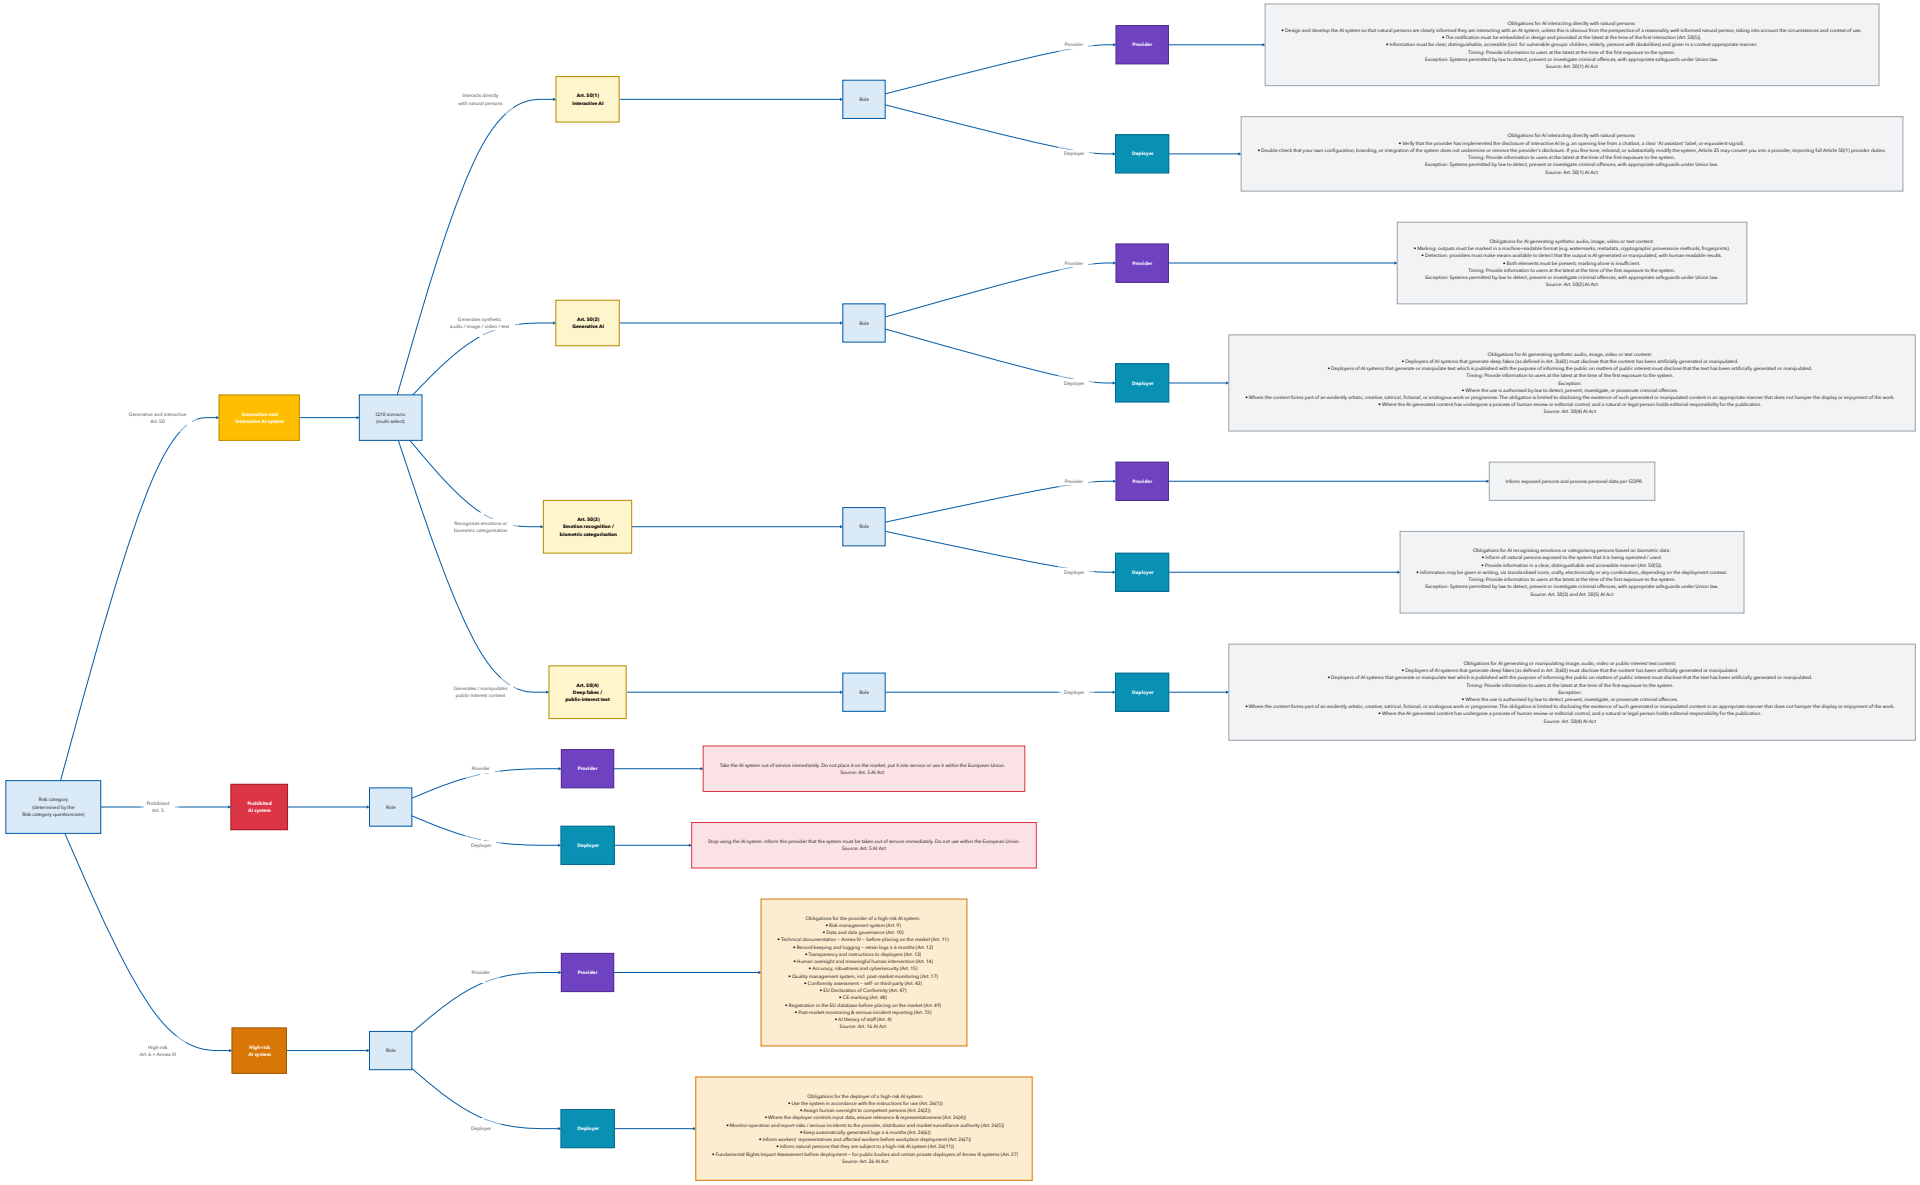

About Algorithm Audit

Algorithm Audit is a European knowledge platform for AI bias testing and normative AI standards. The goals of the NGO are four-fold:

Knowledge platform	Bringing together experts and knowledge to foster the collective learning process on the responsible use of algorithms, see for instance our AI Policy Observatory and position papers
Normative advice commissions	Forming diverse, independent normative advice commissions that advise on ethical issues emerging in real world use cases, resulting over time in algoprudence
Technical tools	Implementing and testing technical tools for bias detection and mitigation, e.g, bias detection tool, synthetic data generation and sociotechnical evaluation of generative AI
Project work	Support for specific questions from public and private sector organisations regarding responsible use of AI

Structural partners of Algorithm Audit

SIDN Fund

The SIDN Fund stands for a strong internet for all. The Fund invests in bold projects with added societal value that contribute to a strong internet, strong internet users, or that focus on the internet's significance for public values and society.

European AI&Society Fund

The European AI&Society Fund supports organisations from entire Europe that shape human and society centered AI policy. The Fund is a collaboration of 14 European and American philantropic organisations.

Building *AI auditing* capacity from a *not-for-profit* perspective

www.algorithmaudit.eu

www.github.com/NGO-Algorithm-Audit

info@algorithmaudit.eu

Parkstraat 22, 2514 JK Den Haag

Algorithm Audit

Stichting Algorithm Audit is registered as a non-profit organisation at the Dutch Chambre of Commerce under license number 83979212