Algorithm Audit

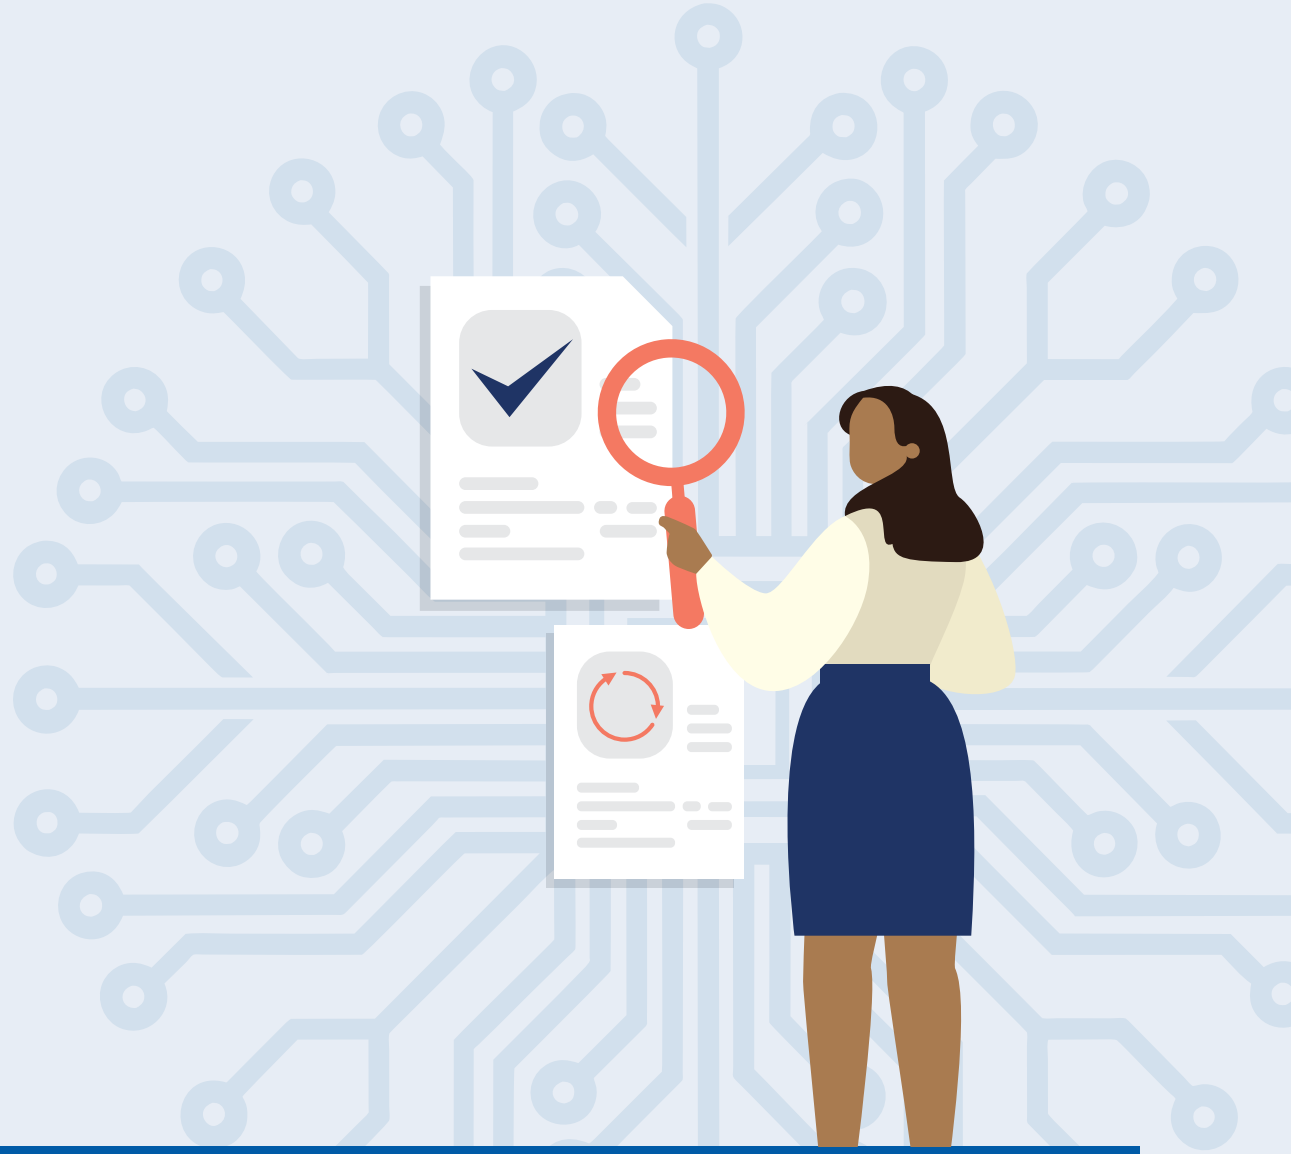

AI AQT documentation

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

March 2026

Table of Contents

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

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 two user-friendly, dynamic questionnaires designed to support the harmonized identification and risk classification of algorithmic systems.

Questionnaire 1 focuses on:

- **AI systems:** As defined in the AI Act.
- **High-impact algorithms:** Systems with severe impact on stakeholders monitored by the Dutch government (irrespective of their status under the AI Act).
- **Solely automated decision-making (sADM), including profiling:** Automated decision-making (ADM) practices as restricted or prohibited under article 22 of the GDPR.

Questionnaire 2 focuses on:

- **Prohibited AI systems:** AI systems as defined in article 5 of the AI Act.
- **High-risk AI systems:** AI systems that require additional control measures as defined in article 6 and Annex III of the AI Act.
- **Transparency requirements:** Additional transparency requirement for certain AI systems as set out in article 50 of the AI Act.
- **General Purpose AI (GPAI):** Requirement for the GPAI model provider as set out in Article 53 of the 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 1. 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 1 and 2. Algorithm Audit has also published supplementary materials on interpreting the AI Act's definition of an AI system[1], as well as on the scope of sADM and the requirement of meaningful human intervention.[2]

The AI AQT is aimed at professionals without a legal background, who can make use of the tool without consulting legal or compliance experts. In the questionnaires, Q1-Q3 and Q8-Q10 focus on the system's logic and output, followed by Q4-Q7, which concern the process in which the algorithmic system is used. To ensure user-friendliness, legal definitions are translated into natural 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

[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).

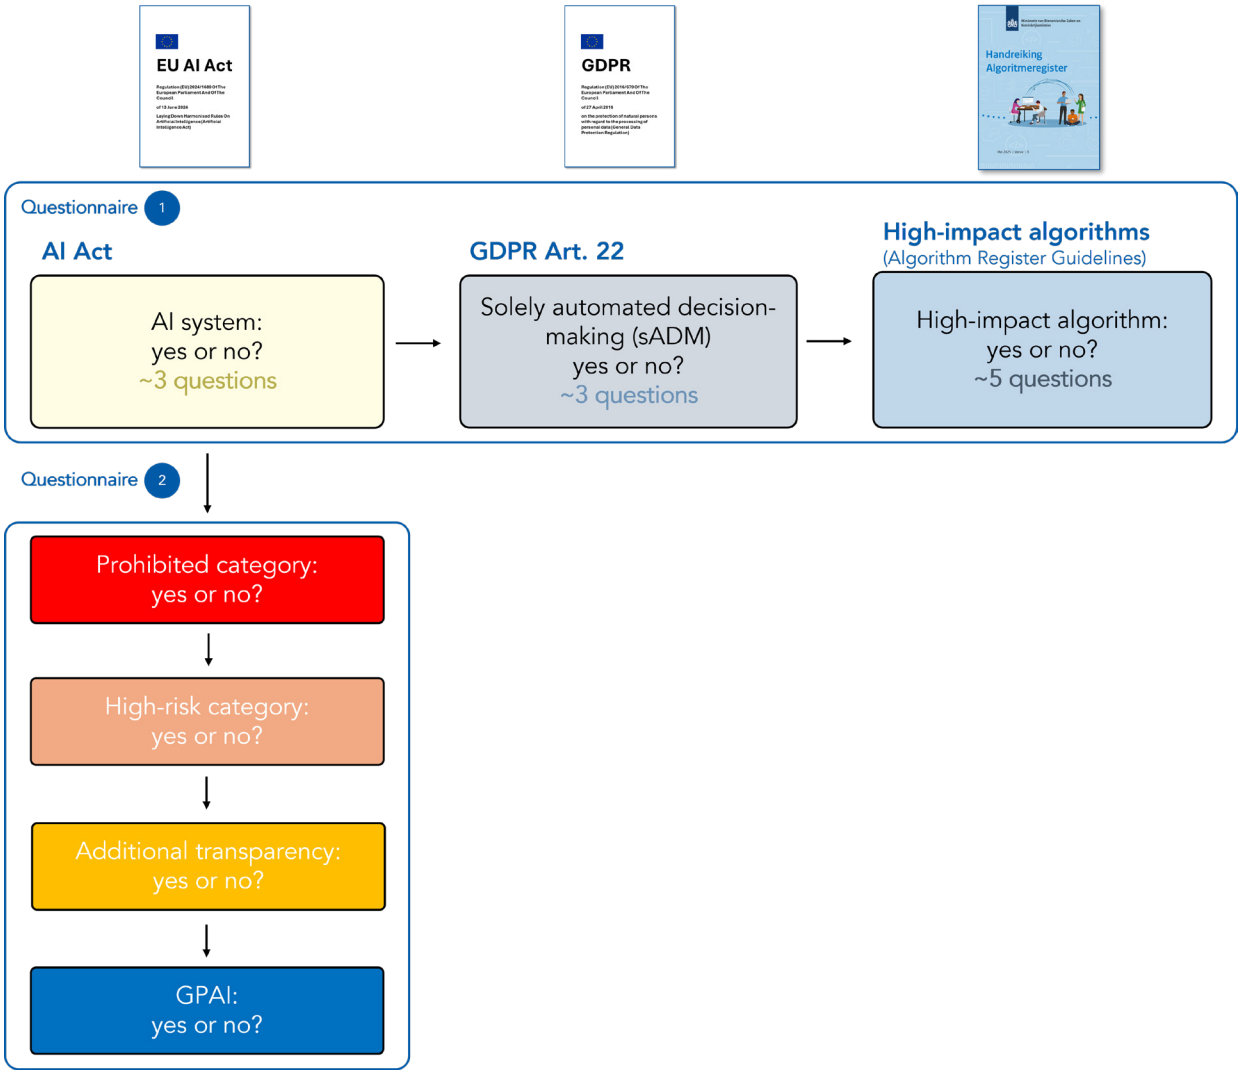

Figure 1 - Overview of the flow of dynamic questionnaires in the AI AQT. Responses to earlier questions are taken into to account in order to minimize the number of questions presented to the user.

is out of scope. The tool is intended for guiding further compliance actions and does not substitute proper legal assessment or consultation.

This paper details the considerations and design choices made during development of the AI AQT. It first explains the components of Questionnaire 1 used to determine whether an algorithmic system qualifies as an AI system under the AI Act (section 2). It then introduces the concept of high-impact algorithms and elaborates how the tool identifies such systems (section 3). In conclusion, it makes clear the questions used to identify sADM following article 22 of the GDPR (section 4). The document concludes by describing how Questionnaire 2 works to qualify AI systems that require additional control measures (section 5).

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.

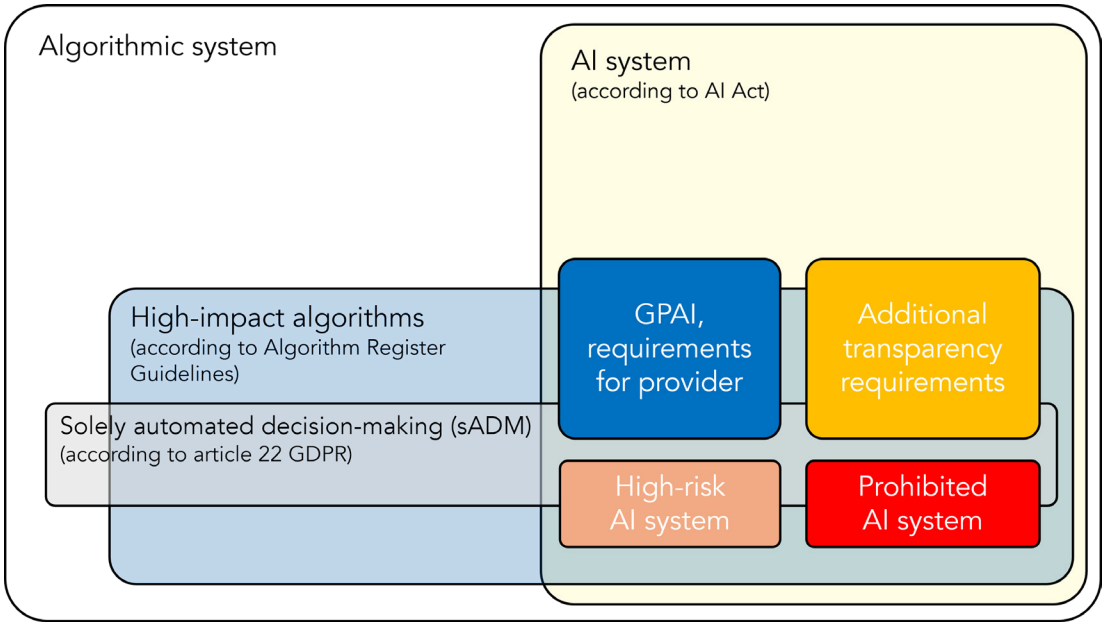

Figure 2 - Overview of the type of algorithmic systems identified by the AI AQT.

2. Questionnaire 1: AI system

The definition of an AI system 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 this 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'.

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

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 an AI system and sADM.

This section begins with an explanation of what is meant with inference (2.1) and autonomy (2.2). Guidelines provided by the EC (hereafter: 'guidelines') on the interpretation of the definition of an AI system are incorporated in this analysis. [4]. Next this section discusses how these concepts are incorporated into the AI AQT Questionnaire 1 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 of an AI system.[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 definition of AI systems. 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 definition of AI in the AI Act (hereafter: '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 1 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 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'.

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.

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 distinguish AI systems from other algorithmic systems.

Q1 – What is the outcome of the application?

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 dashboards do not transcend *"the elementary processing of data by enabling learning, reasoning or modelling"*[8] as an AI

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

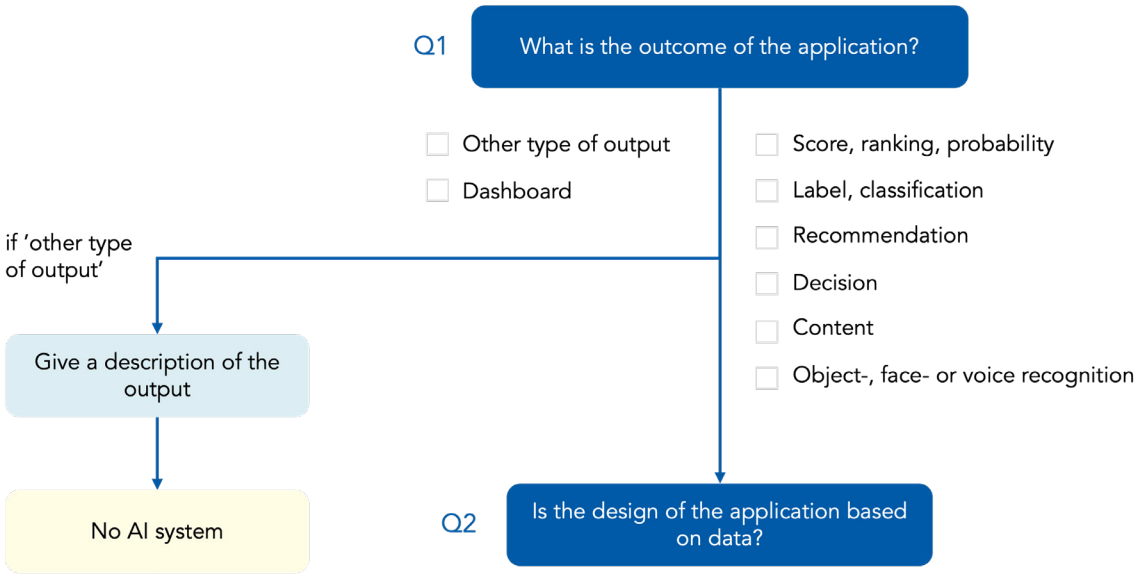

Figure 3 - Q1 asks what type of output the algorithmic application generates.

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.

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.

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

While 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 1: High-impact algorithms and 4. Questionnaire 1: Solely automated decision-making, 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-

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

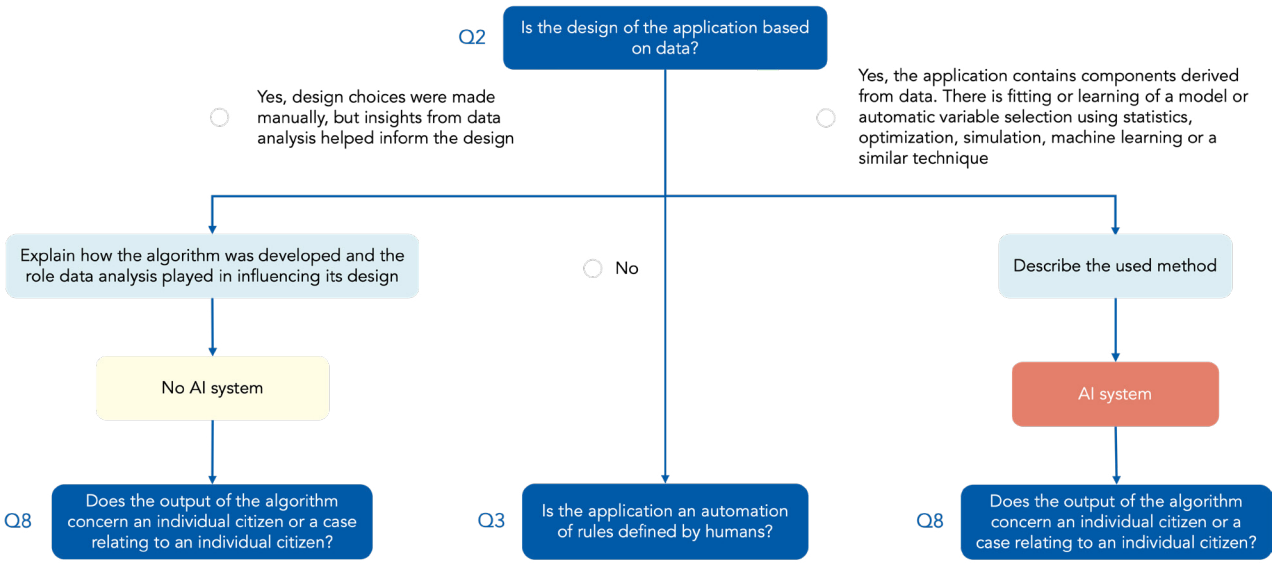

Figure 4 - Q2 is about whether the design of the application is based on data.

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 1: High-impact algorithms and 4. Questionnaire 1: Solely automated decision-making, starting with Q8.

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.

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

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

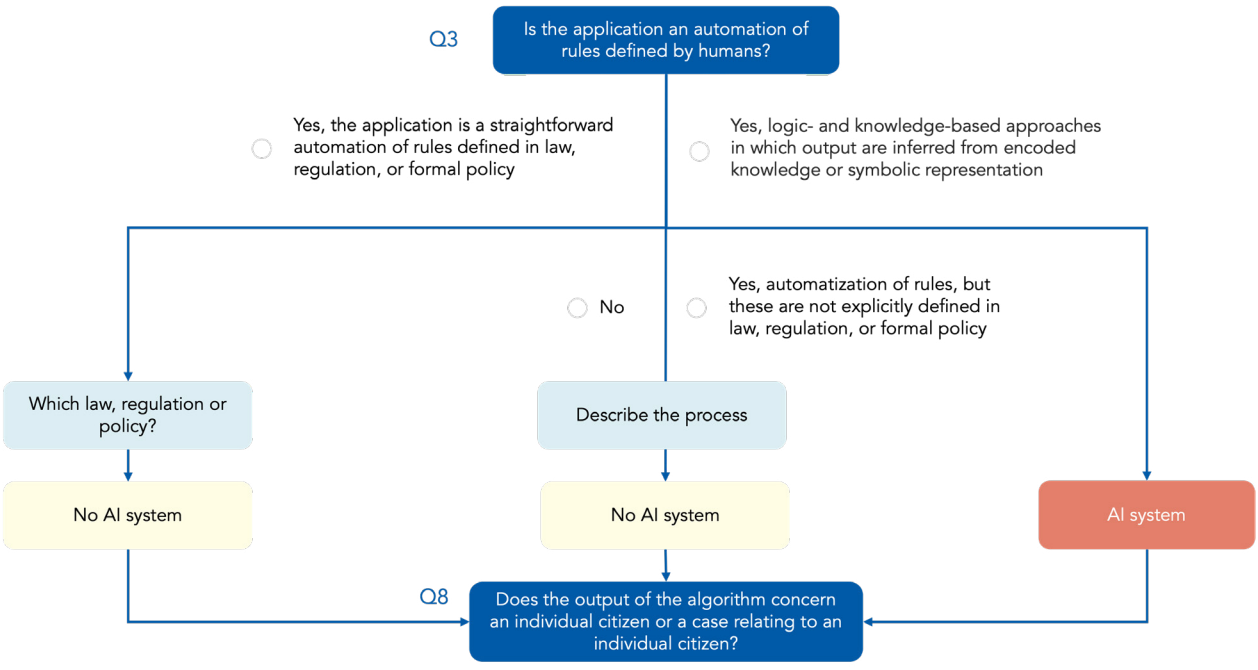

Figure 5 - Q3 examines the extent to which there is human involvement in the creation of rules used in an algorithm.

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 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 1, 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. Thereafter, users are asked whether they want to continue with the rest of Questionnaire 1, which deals with high-impact algorithms and sADM, starting with Q8.

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. After Q3, users are asked whether they want to continue with the rest of Questionnaire 1, which deals with high-impact algorithms and sADM, starting with Q8.

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

Flowchart – AI system (AI Act)

This is a simplified representation of questions asked in the AI AQT. This schematic representation shows the logic required to determine whether there is an *AI system* according to 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.

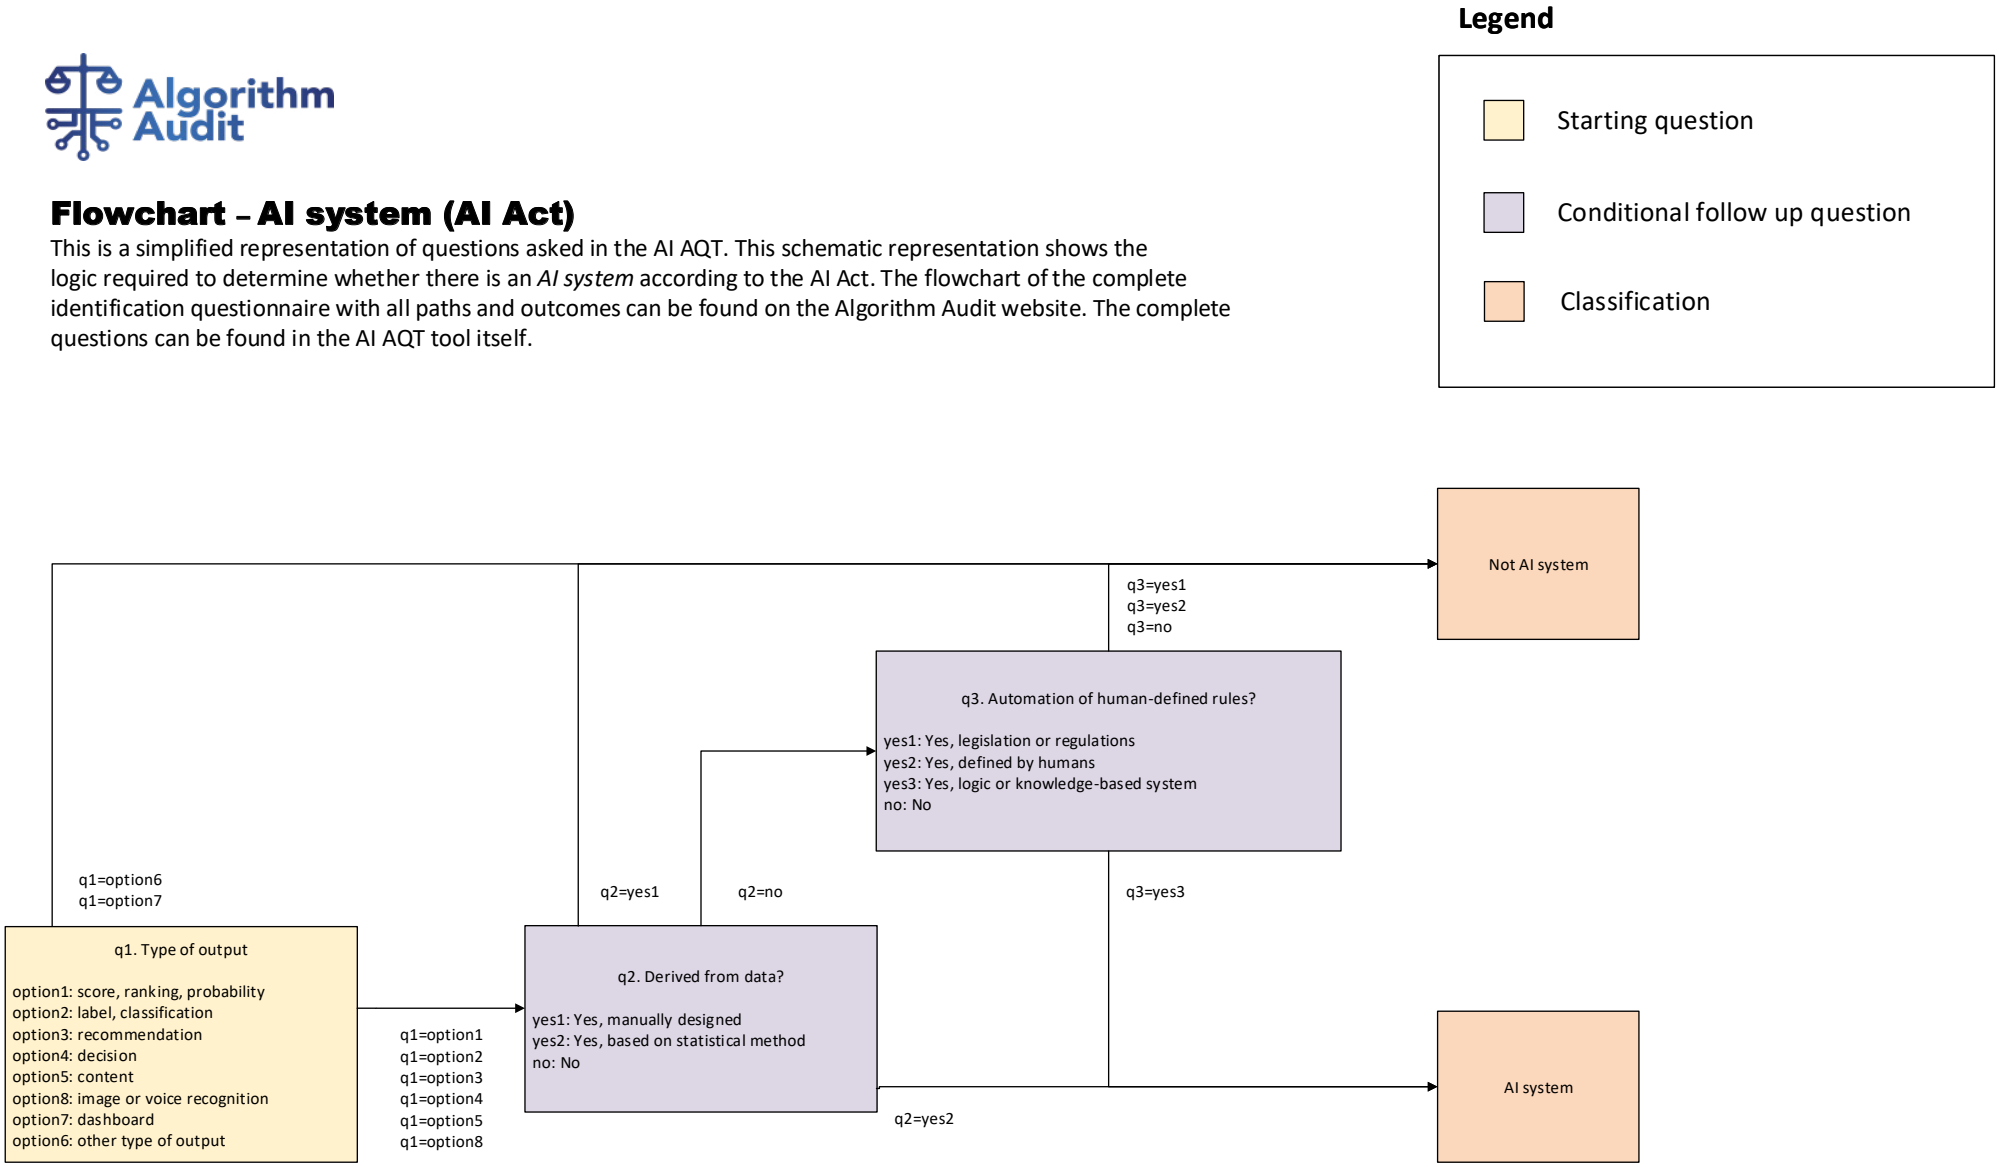

3. Questionnaire 1: High-impact algorithms

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"*.[11] 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:[12]

> **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 6.[13]

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?

[11] 'Aandacht voor algoritmes', The Netherlands Court of Auditors (2021).
[12] Algorithm Register Guideline of the Ministry of the Interior and Kingdom Relations (2023).
[13] Style adapted from Algorithm Register Guideline of the Ministry of the Interior and Kingdom Relations (2023). Supra note 15.

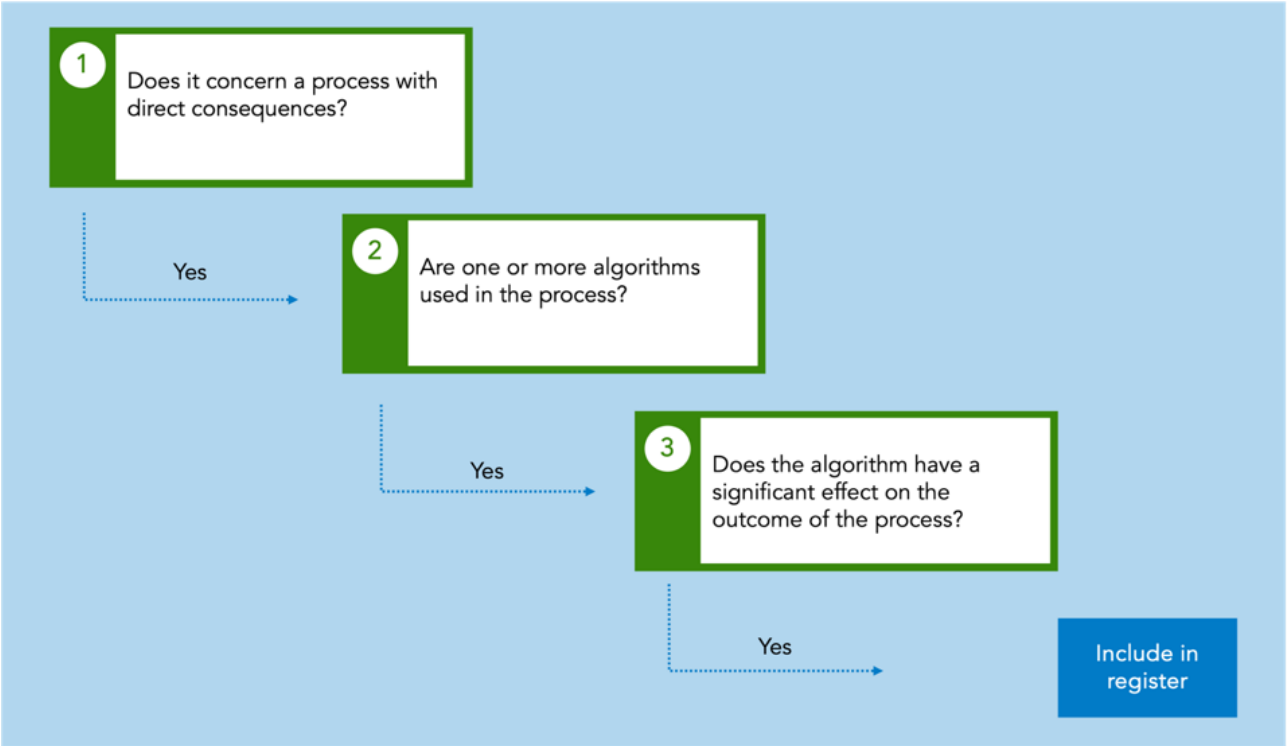

Figure 6 - Questions from Algorithm Register Guidelines that can be used to determine whether an algorithm is a 'high-impact algorithm'.

This section discusses how the first and last questions in Figure 6 – about the process in which the algorithm is used – are incorporated into the AI AQT as Questions 4-7 (Q4-7). Assuming at least one algorithm is involved in a given decision-making process, it can be deemed the second question in Figure 6 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' (3.1) and a 'significant effect' (3.2).

NOTE: High-impact algorithms must be published in the Dutch national Algorithm Register[14], unless there is a ground for exception.[15]

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

[15] Grounds for exception as stated in the Guideline are: "legal grounds for exception as specified in the Open Government Act (Woo) and the Public Health 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.[16] 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.[17] 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.[18] 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.

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

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

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

3.1 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":[19]

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 (Q4). The notion of a decision here is to be interpreted broadly: not only formal decisions, as defined in Dutch public administration Law (Awb art.1:3), 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 (Q5).

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 6 (Q6).

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.

3.2 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":[20]

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*

[19] Supra note 15.
[20] Supra note 15.

that (partly) determines the next step in the 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. [21] 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 Q7.

NOTE: The logic for the flow chart for the answer options in Q3 are different when assessing high-impact algorithms than when assessing an AI system. See 2. Questionnaire 1: AI system.

Q1 – What is the outcome of the application?

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 7.

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 3.2.II). If one of these options is selected the user is brought to Q2.

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

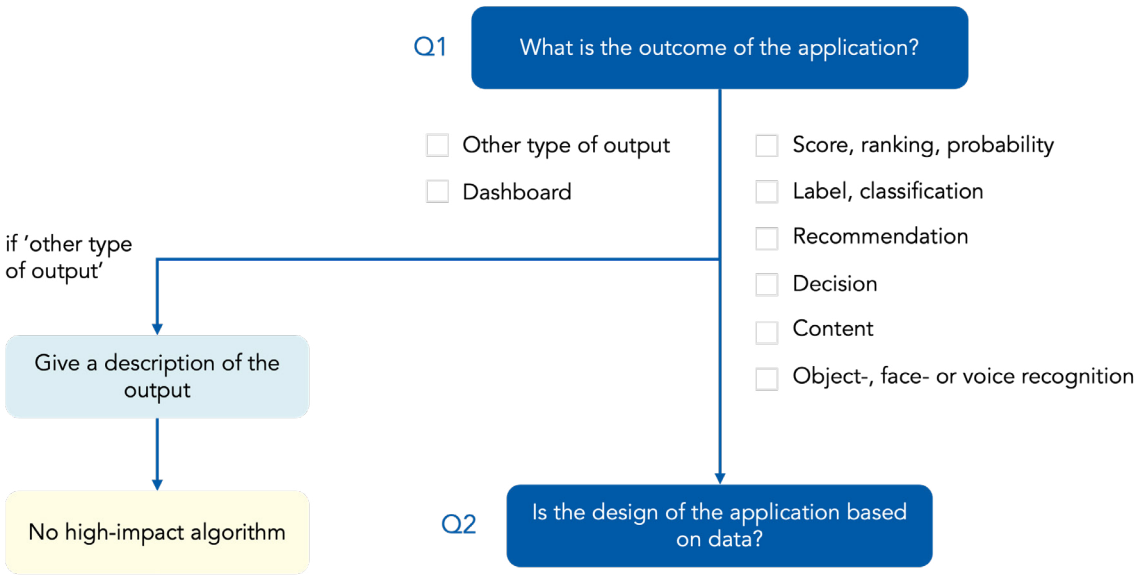

Figure 7 - Q1 concludes that the application does not qualify as a high-impact algorithm if a 'Dashboard' or 'Other type of output' is at stake.

Since a dashboard, on its 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.

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 3.2.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 to determine the consequences of algorithmic system.

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.

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

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 3.2 Significant effect on the outcome of the process and Figure 9.

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 3.2 Significant effect on the outcome of the process) and the user is asked whether they wish to proceed to Q4.

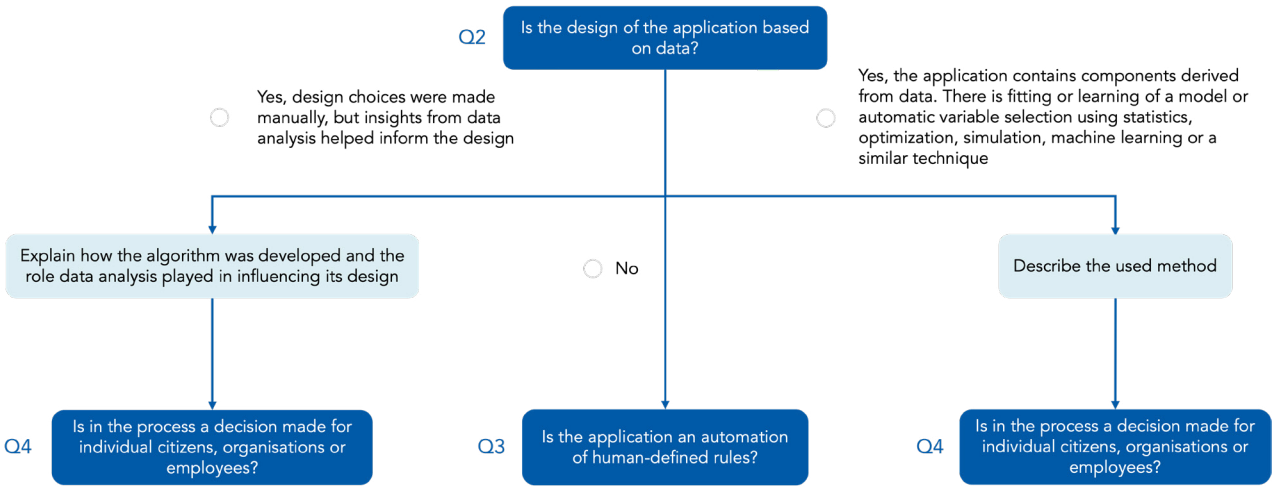

Figure 8 - Q2 distinguishes whether more information is needed about the design of the algorithm (Q3), or that the user can continue to the next part of Questionnaire 1 (Q4).

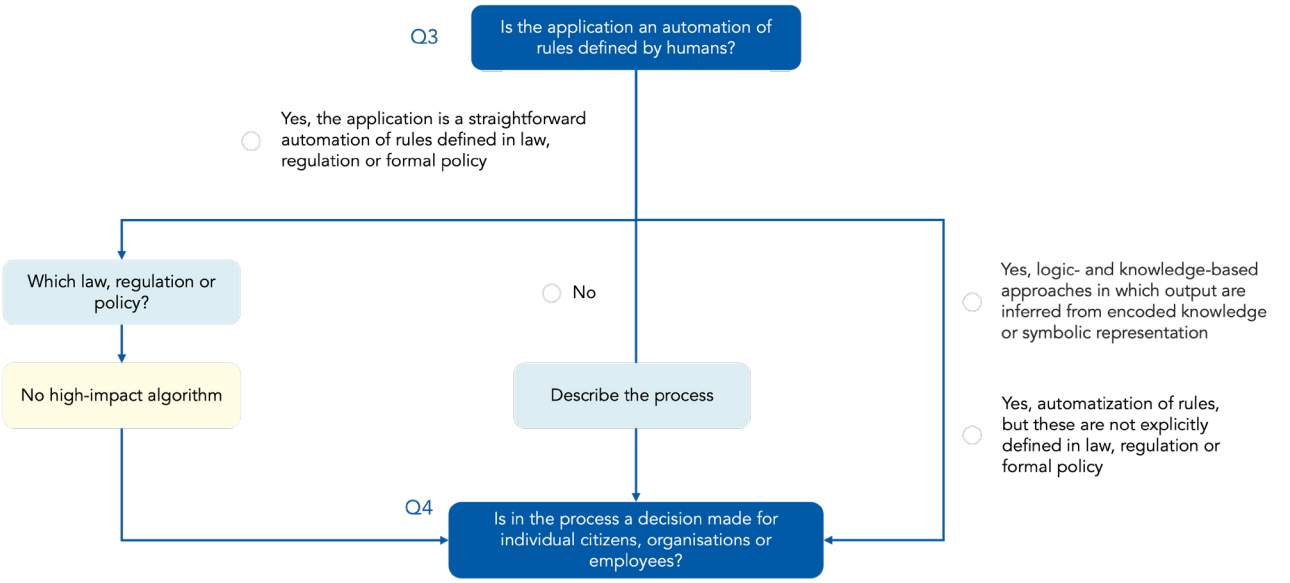

Figure 9 - In case of 1-to-1 automation of defined rules in law, regulation or formal policy, Q3 concludes that the application is no high-impact algorithm. Other scenarios are forwarded to Q4.

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 1: 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.

Q4 – Is in the process a decision made for individual citizens or civil servants?

To identify an essential characteristic of a high-impact algorithm (see section 3.1), it is first assessed whether a decision is made in the process the algorithm is involved. See Figure 10.

In case a decision is made, the user is forwarded to Q5. In case a decision is not made, the user is forwarded to Q6.

To assist users, the following remark is provided:

Consider prioritizing the follow-up to a citizen's question or request, whether or not to ask a citizen to provide additional information, whether or not to select someone for a check or inspection, whether or not a person is eligible for services or facilities, etc.

NOTE: A decision is much broader than a formal decision as defined in Dutch National Administrative Law.

Q4 relevant for qualifying high-impact and sADM

Q4 also supports the qualification of an algorithmic as sADM. This is further explained in 4. Questionnaire 1: Solely automated decision-making.

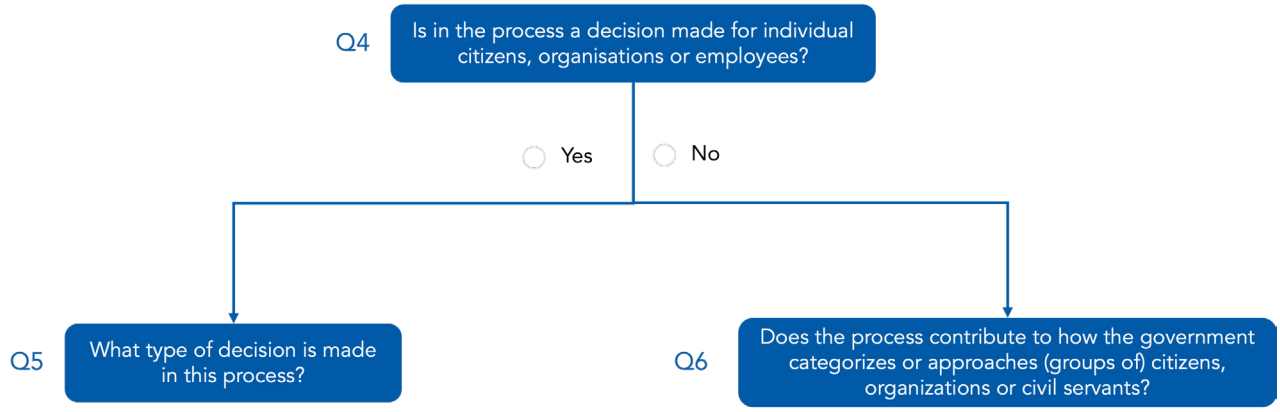

Figure 10 - Q4 examines whether a decision is made in the process in which an algorithm is involved.

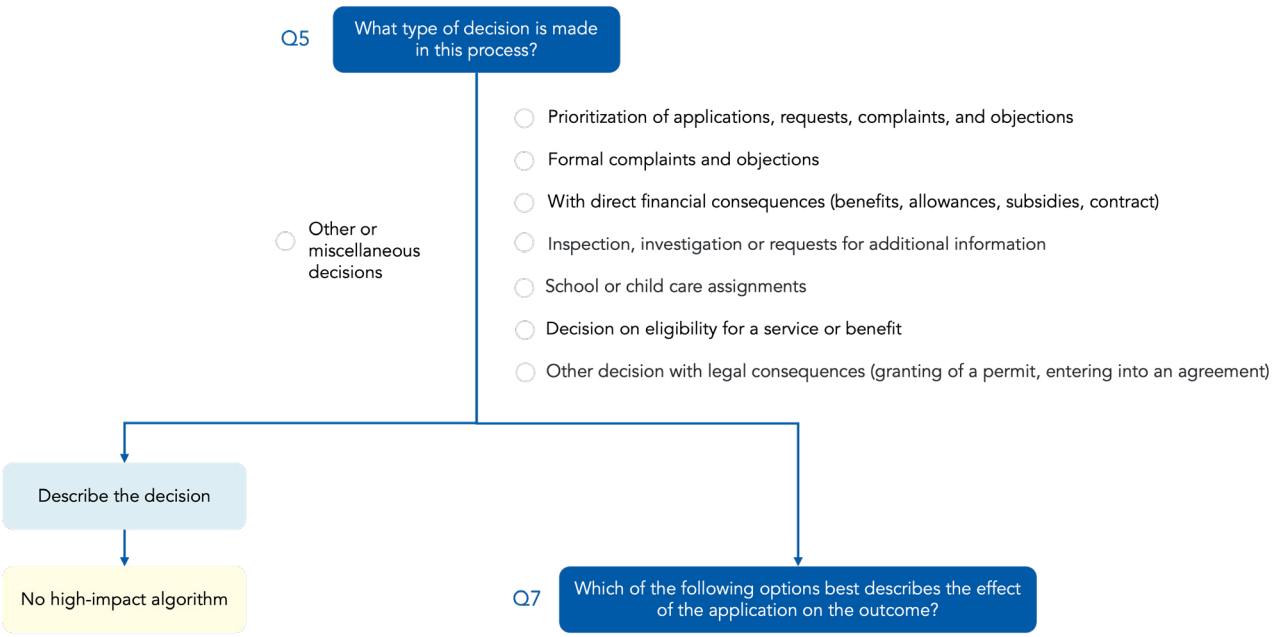

Figure 11 - Q5 asks what kind of decision is made in the process.

Q5 – 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 3.1 Direct consequences and Figure 11.

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 Q7. Note that only one option can be selected.

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: Q5 also supports the qualification of a system as sADM. This is further explained in 4. Questionnaire 1: Solely automated decision-making.

Q6 – 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 Q4, Q6 examines whether the process contributes to how the government categorises or approaches groups of citizens, organisations or civil servants. See Figure 9.. This is the second aspect of checking whether a direct consequence might follow from applying the algorithm (see section 3.1).

If the process does not contribute to how the government categorises or approaches groups of citizens, organisations or civil servants, 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 Q7. The user will also be redirected to Q7 if Q6 is answered with 'Yes'.

NOTE: Q6 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 are approached by governmental actors.

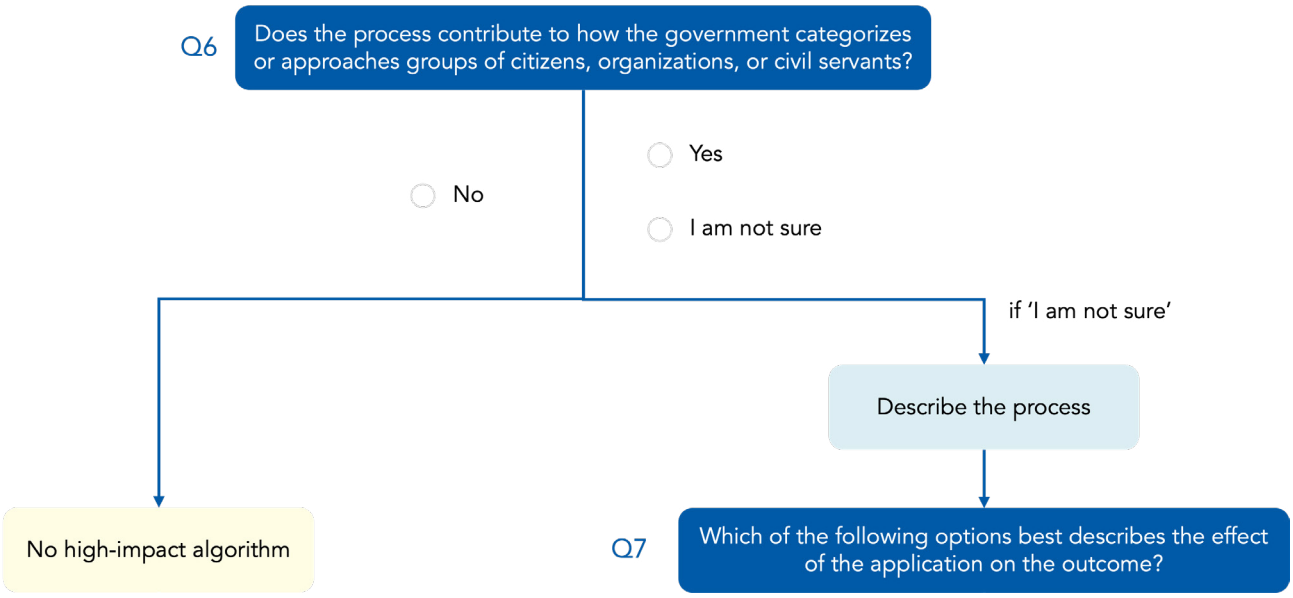

Figure 12 - Q6 examines whether the process contributes to how (groups of) citizens, organisations or civil servants are categorised or approached by the government.

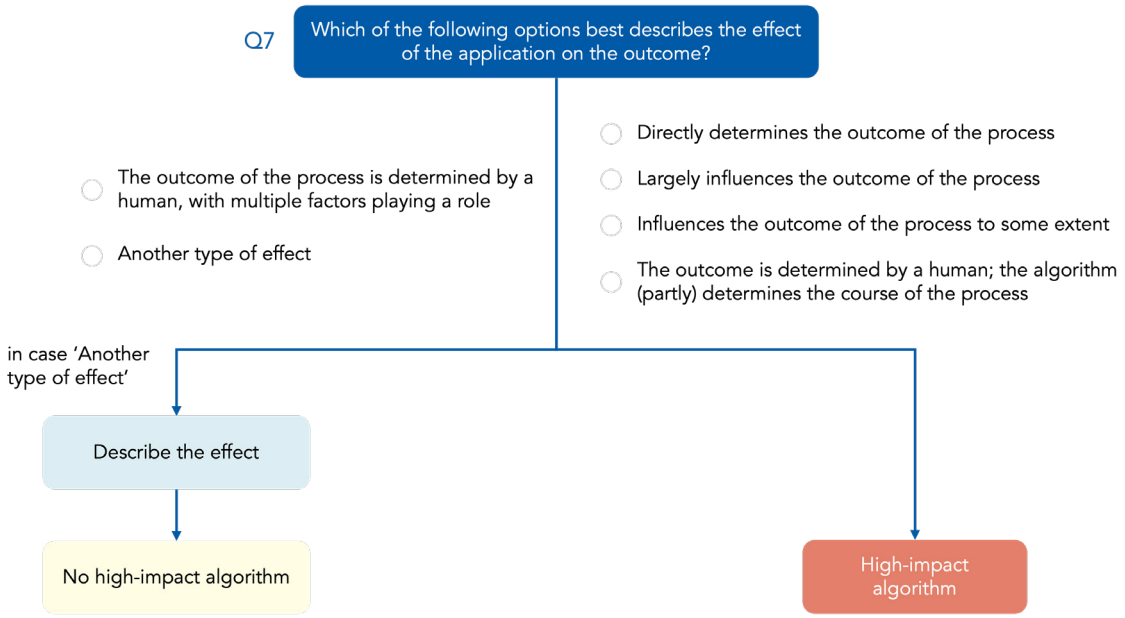

Figure 13 - Q7 examines what the effect of the outcome of the algorithm is on the outcome of the process.

Q7 – Which of the following options best describes the effect of the application on the outcome?

After it is established whether a decision is taken (Q4 and Q5) or the process alters how the government categorizes or approaches data subject (Q6), it must be determined whether the algorithm significantly effects the outcome of the process. If both conditions are met, it can be concluded algorithm is a high-impact algorithm. See Figure 13.

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, largely influences or influences the outcome to some extent, it is considered to have a significant effect and therefore qualifies as a high-impact algorithm. This also applies where the outcome is formally determined by a human, but the algorithm (partly) determines the course of the process.

Where the outcome of the process is determined by a human and multiple factors – other than 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: Q7 also supports the qualification of a system as sADM. This is further explained in 4. Questionnaire 1: Solely automated decision-making.

Flowchart identification high-impact algorithm

Algorithm Audit

Flowchart – High-impact algorithm (Algorithm Register Guidelines)

This is a simplified representation of questions asked in the AI AQT. This schematic representation shows the logic required to determine whether there is a *high-impact algorithm* according to the Algorithm Registry Guidance Document. 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.

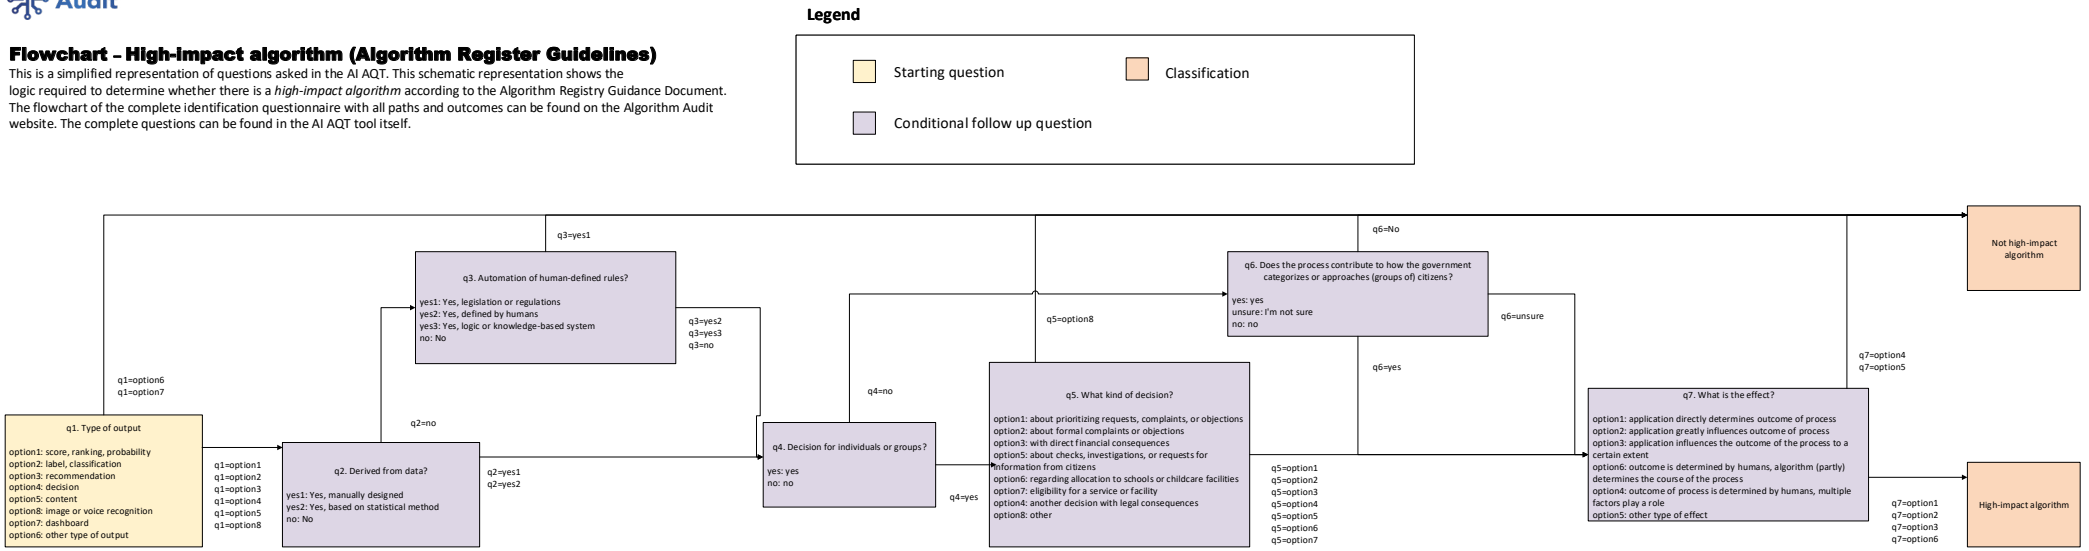

4. Questionnaire 1: Solely automated decision-making

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"* [22] 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"*[23] published by the Dutch Data Protection Authority (DPA), *"Advice on automated selection techniques"* by legal experts at Pels Rijcken[24] and relevant legal-scientific literature[25,26]. Algorithm Audit has also published supplementary materials providing a step-by-step guide how decision based solely on automated processing can be prevented.[27]

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 (4.1). 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. As such, the same questions can be used to qualify high-impact algorithms and sADM (i.e., Q5 and Q7). Hence, section 4.1. explicitly links to concepts explained in 3. Questionnaire 1: High-impact algorithms. Additionally, based on the analysis in this section, four additional questions are included in the AI AQT (Q8-Q10 and Q5.1) to identify sADM.

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 3.2 Significant effect on the outcome of the process).

4.1 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.

I. A decision is made

The notion of a decision must be interpreted broadly: not only formal decisions, as defined in Dutch public

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

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

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

[25] 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).

[26] 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).

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

administration Law (Awb art.1:3) 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.[28]

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:[29]

- **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"*.

Significant consequences for stakeholders might also arise if the outcomes of a risk profiling algorithm are shared internally or externally, or are stored long-term, 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 oversight of the decision is meaningful, rather than just a token gesture. It should be carried out by*

[28] ECLI:EU:C:2023:957, case C-634/21, Court of Justice of the European Union (2023).
[29] Supra note 22-26.

someone who has the authority and competence to change the decision. As part of the analysis, they should consider all the relevant data".[30]

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 determines or largely influences 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[31], 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"*.[32]

Q1 – What is the outcome of the application?

The type of output generated by an algorithmic system gives an indication of whether the system qualifies as sADM. See Figure 14.

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 4.1). If one of these options is selected the user is brought to Q2.

[30] Supra note 22.
[31] Supra note 23 and 22.
[32] Supra note 28.

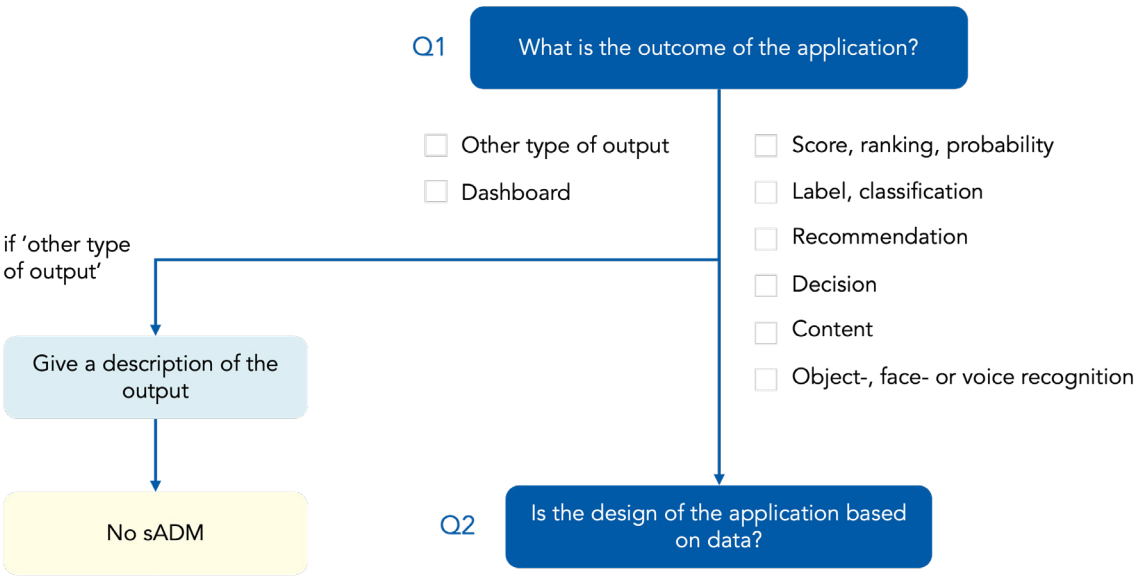

Figure 14 - Q1 concludes that the application does not qualify as sADM if a 'Dashboard' or 'Other type of output' is selected.

Since a dashboard, on its own, only provides data visualization, there is no decision, nor 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.

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.

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

The only relevant difference in how Q3 captures the concepts of high-impact algorithms and sADM is that the form one-to-one automation cannot qualify as a high-impact algorithm but may still constitute sADM. Before posing Q4-Q7, users are first presented with an additional set of questions exclusively pertaining to sADM – Q8-Q10. Additionally, as explained below, Q9 and Q10 cover a special scenario where, even if unused, simply generating and storing or sharing process outcomes can qualify as sADM, following case law. See Figure 15.

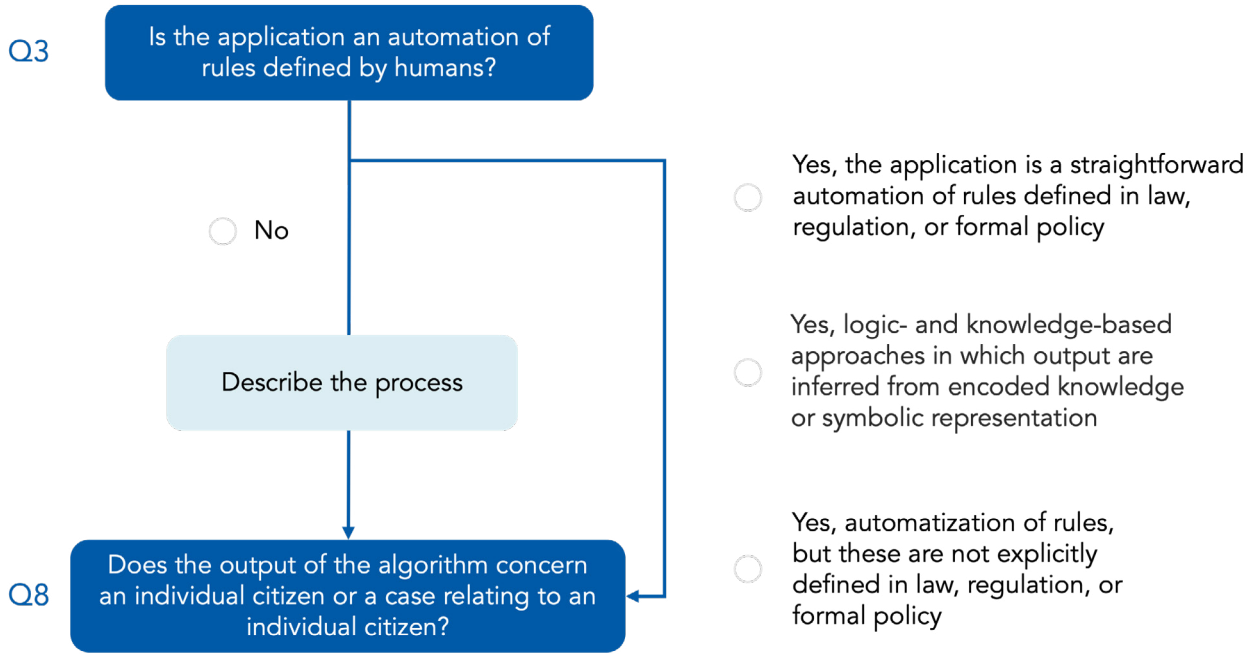

Figure 15 - All Q3 responses route users to Q8 for sADM identification.

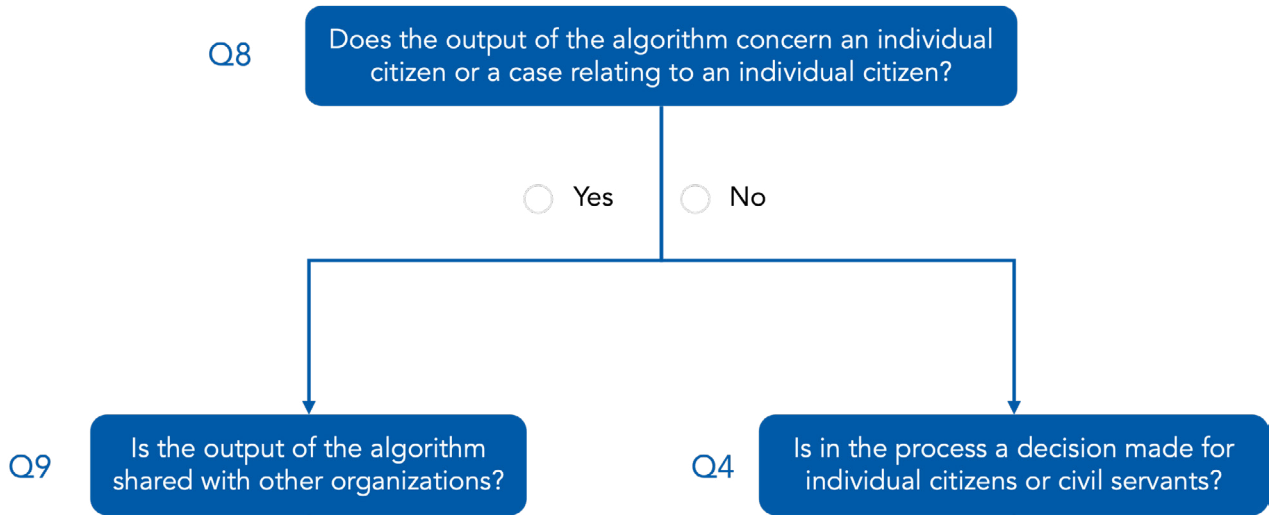

Figure 16 - Q8 determines whether the output of the algorithm relates to an individual citizen or to a case involving an individual citizen and should be regarded as personal data or not.

Q8 – Does the output of the algorithm concern an individual citizen or a case relating to an individual citizen?

Since article 22 explicitly refers to decisions pertaining to individuals, Q8 inquiries about the output of the decision-making process. If the output is not directly about or cannot be traced to an individual, the process may not constitute sADM. Consider the example of a neighbourhood profiling algorithm used to allocate police presence. Downstream, this may well impact an individual, but the output of the system that led to the decision with significant effect cannot be traced to an individual. Hence there are no grounds for prohibition under article 22. This is a notable difference between the scope of high-impact algorithms and sADM. The former can apply to groups, the latter only to individuals. Hence, Q8 is needed in parallel to Q4. See Figure 16.

In case the output of the algorithm concerns an individual citizen or a case relating to an individual citizen, the user is directed to Q9 to assess how the output of the algorithm is used.

Even if this is not the case, the user is still redirected to Q4. This follows the "better safe than sorry" principle of the design. In practice, users are expected to also answer Q4 with 'No.' Since Q4 can also capture whether a decision pertaining to an individual is made, the user is not immediately presented the conclusion that sADM is not at hand.

To assist users, the following remark is provided:

Examples of output relating to an individual citizen include:

> An assessment of an individual characteristic.

Examples of cases relating to an individual citizen include:

> Matters linked to an individual, such as a transaction or an application.

Examples of information that does not relate to individuals include:

> Output relating to groups, where individuals are not assigned an outcome separately from the group
> Output relating to physical objects that are not linked to an individual
> Output relating to sectors or neighbourhoods
> Output relating to financial policy and the effects of policy.

When considering individuals, this also includes businesses for which the owner is personally liable (e.g., sole proprietorships, one-person businesses, general partnerships, and professional partnerships).

Q9 – 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 4.1).[33] Q9 therefore assesses whether the output of the algorithm is shared with other organizations. See Figure 17.

When the output of the algorithm is stored for longer than the duration of the primary process for which it is used, users are presented with a warning (at the end of the questionnaire) that storing 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 Q10.

=Q10 – 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 4.3). Q10 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 19.

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 Q4 to continue the questionnaire.

[33] Additional explanation of the Shufa ruling can be found in supra note 27.

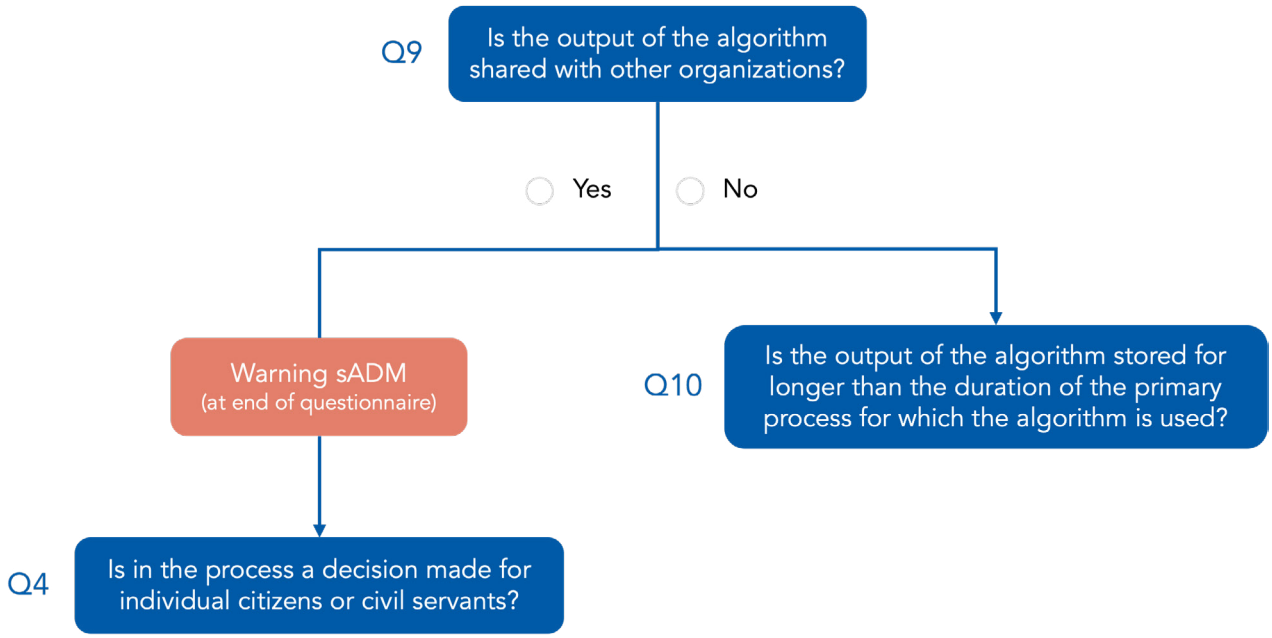

Figure 17 - Q9 assesses whether the output of the algorithm is shared with other organizations.

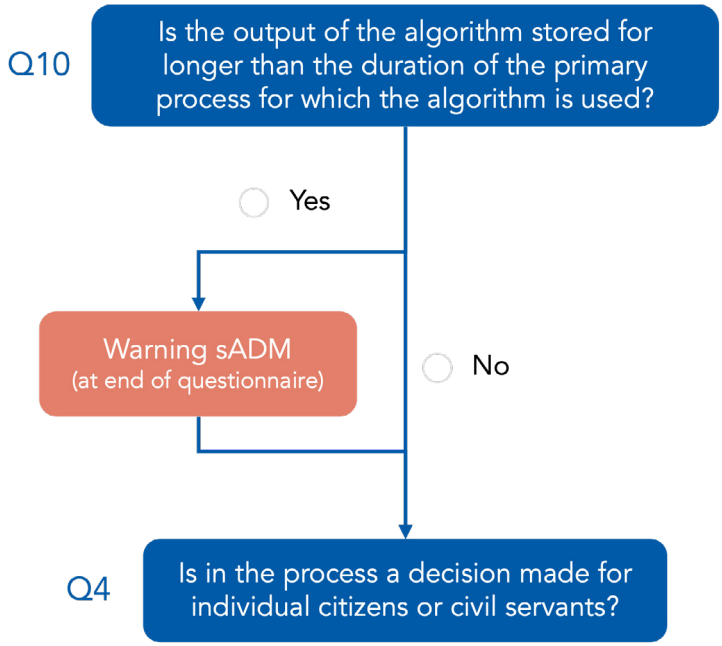

Figure 18 - Q10 examines if the output of the algorithm is stored for longer than the duration of the primary process for which the algorithm is used.

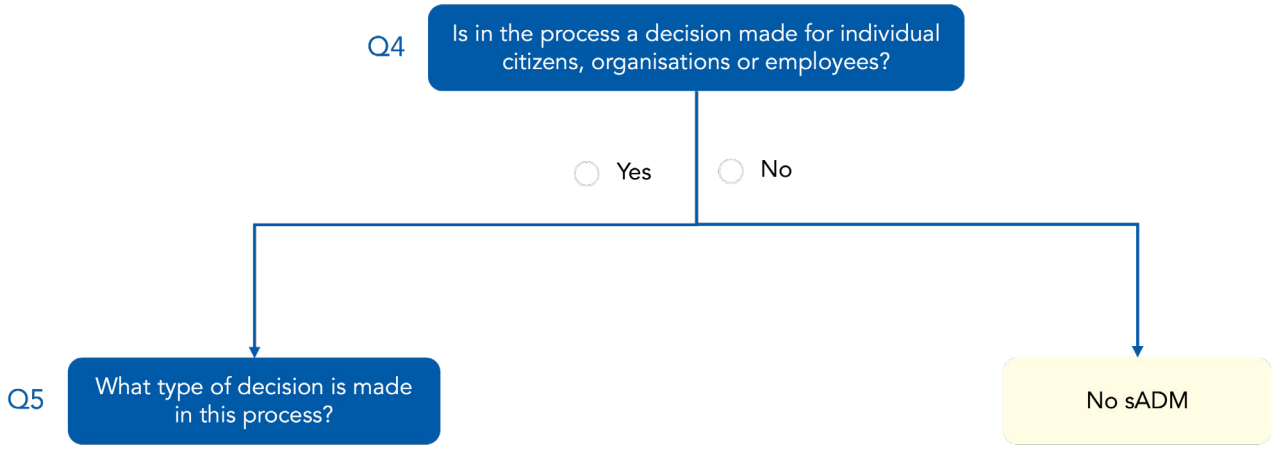

Figure 19 - Whether a decision is made is a key aspect of sADM identification.

Q4 – Is in the process a decision made for individual citizens or civil servants?

Q4 assesses whether a decision is made in the process the algorithm is involved in. Different than for identification of a high-impact algorithm, by answering Q4 with 'no' it is concluded that no sADM takes place. When answering 'yes' users are taken to Q5. See Figure 19.

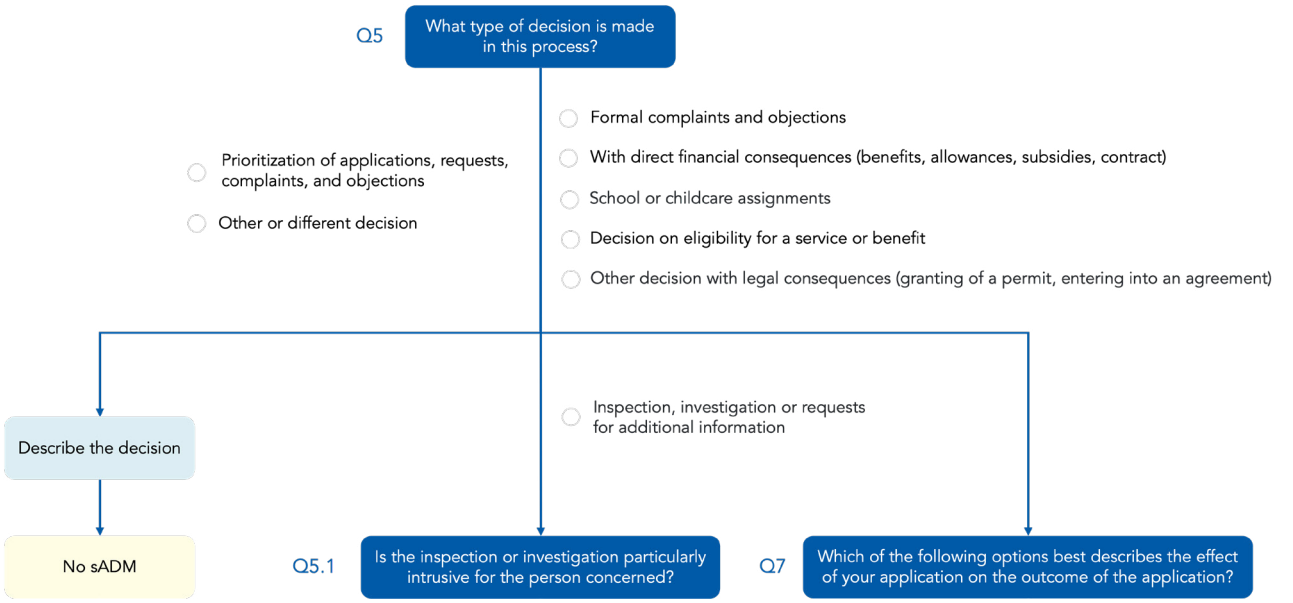

Figure 20 - For sADM identification, a separate answer option is included to identify secondary effect of a decision made in the process in which the application is used.

Q5 – 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 Q5. This serves to assess whether direct consequences follow from the (see section 3.1). See Figure 20.

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 entered, it can be concluded that no sADM occurs. The user is asked for a clarification. This can be reviewed by proper legal experts.

In case Q5 is answered with 'Inspection, investigation, or requests for additional information' an additional question (Q5.1) is asked about the secondary effects.

For all other answers, users are forwarded to Q7.

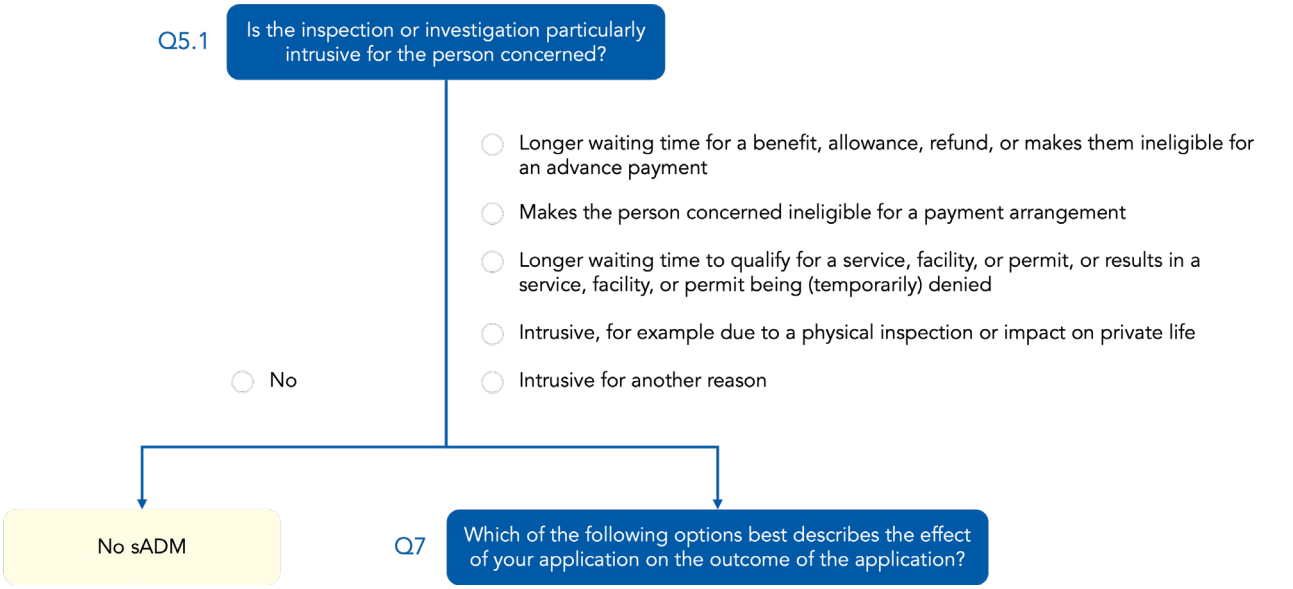

Figure 21 - Q5.1 examines whether the inspection or investigation is particularly intrusive for the person concerned.

Q5.1 – Is the inspection or investigation particularly intrusive for the person concerned?

As mentioned, Q5.1 is shown to the user only if 'Inspection, investigation, or requests for additional information' is selected in Q5. In this context, procedural decisions that do not have direct legal effects can still be considered to have a significant impact (see section 4.1). This is assessed in Q5.1. See Figure 21.

In 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 Q7. 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, users are informed that the process does not constitute sADM.

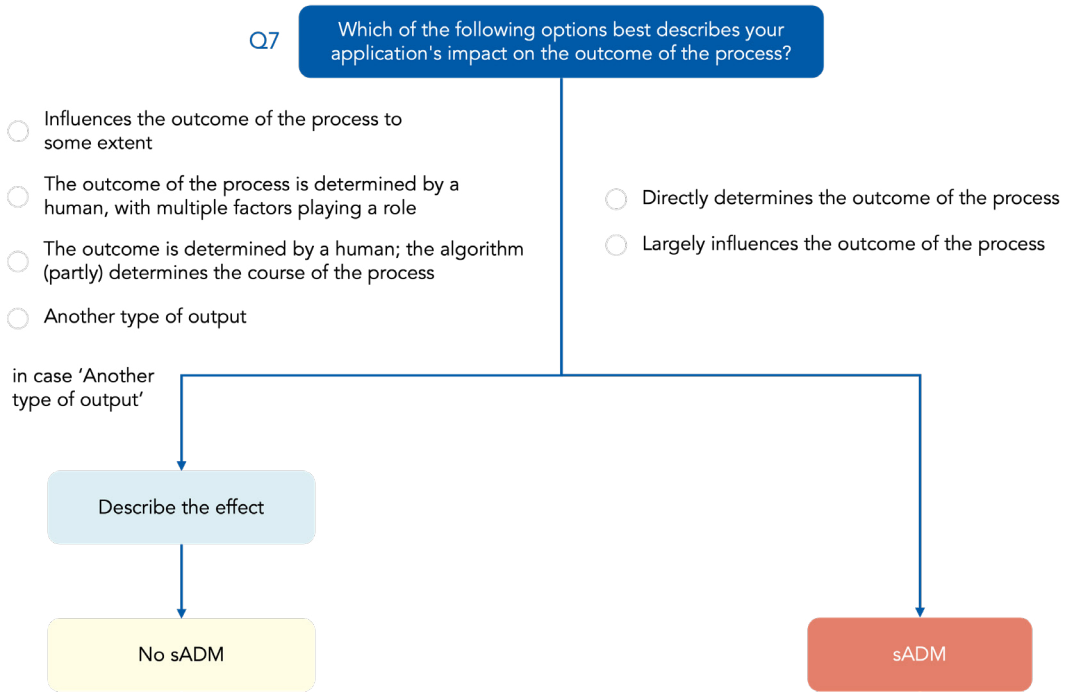

Figure 22 - Q7 examines whether human intervention in the decision-making process is meaningful.

Q7 – Is in the process a decision made for individual citizens or civil servants?

In the context of sADM, Q7 assesses the role of human involvement in the decision-making process and helps determine whether sADM is at hand. See Figure 22.

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.

To assist users, the following remark is provided:

Consider prioritization, follow-up on a question or request from a citizen. For example, whether or not to request additional information from a citizen, whether or not to select for inspection or control, whether or not to (proactively) offer a specific provision in the context of social assistance, etc.

NOTE: A decision is much broader than a formal decision as defined in Dutch National Administrative Law.

Flowchart identification high-impact algorithm

Algorithm Audit

Flowchart – Solely automated decision-making (Art. 22 GDPR)

This is a simplified representation of questions asked in the AI AQT. This schematic representation shows the logic required to determine whether there is *solely automated decision-making* (sADM) according to Article 22 of the GDPR. 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.

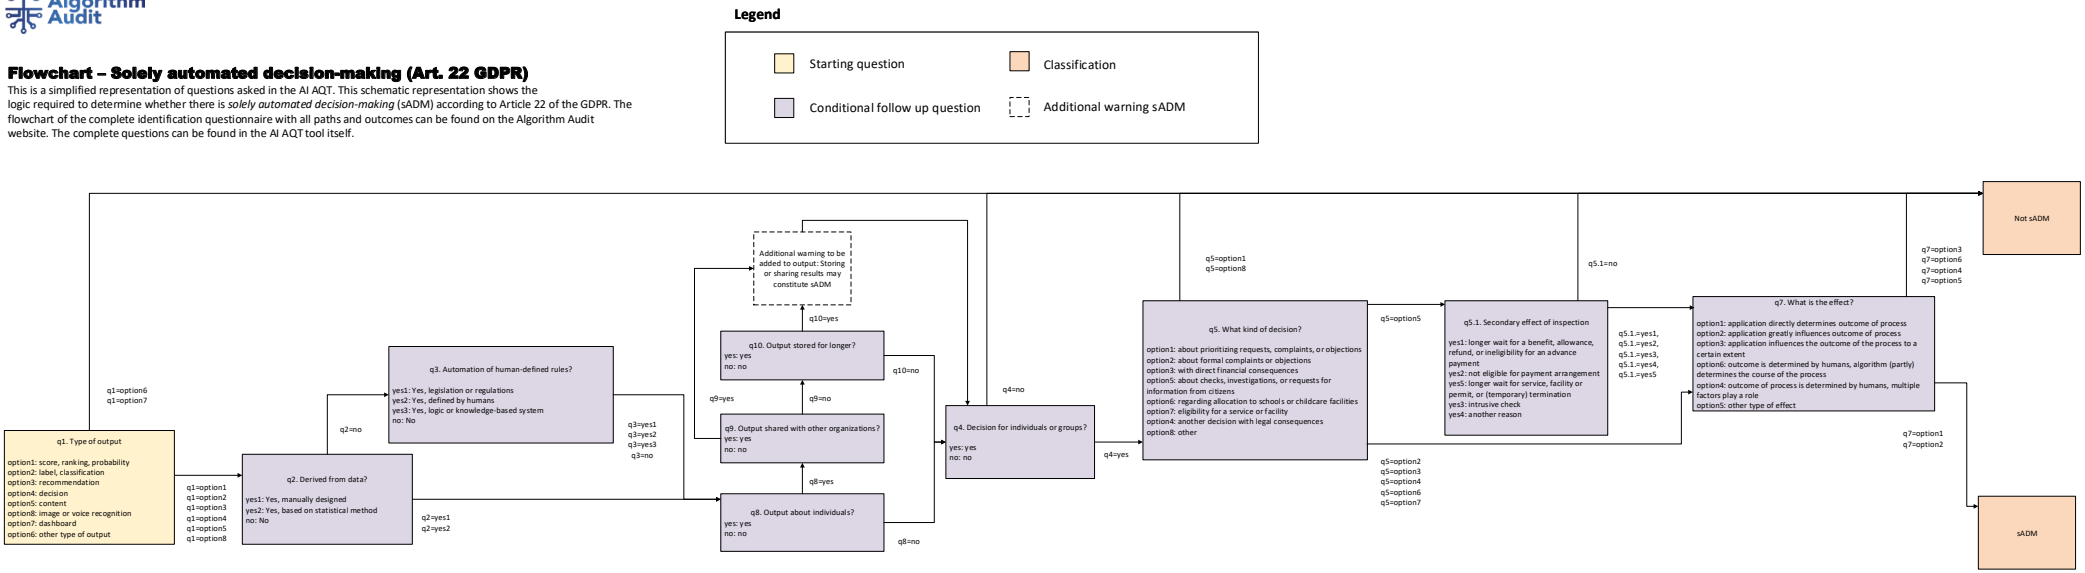

5. Questionnaire 2: Categories of AI systems

The AI Act of the European Union (EU) follows a risk-based approach. Depending on the risk category an AI system belongs to, risk mitigation measures apply. Some AI practices are outright prohibited in the EU (as listed in article 5). For high-risk AI systems (as defined in article 6 and Annex III), requirements apply to both providers and deployers. In addition, transparency obligations apply to providers and deployers of certain AI systems (as mentioned in article 50), and there are specific obligations for providers of general-purpose AI (GPAI) models (as defined in article 53). In Questionnaire 2 of AI AQT, these obligations are summarized in a dynamic questionnaire to determine whether an AI system falls into one of the following categories:

> **Prohibited AI systems**: AI systems as defined in article 5 of the AI Act.
> **High-risk AI systems:** AI systems that require additional control measures as defined in article 6 and Annex III of the AI Act.
> **Transparency requirements:** Additional transparency requirement for certain AI systems as laid down in article 50 of the AI Act.
> **General Purpose AI (GPAI):** Requirement for the GPAI model provider as set out in Article 53 of the AI Act.

NOTE The dynamic questionnaire first determines whether the AI system falls within a high-risk category (if any) listed in Annex III. It then assesses whether the system involved a prohibited practice under Article 5, whether transparency obligations apply under article 50, or whether it qualifies as a GPAI model under article 53. Thereafter, it checks whether any exception under article 2 applies. In the current version of AI AQT, high-risk AI systems intended to be used as safety components of a products covered by EU harmonisation legislation listed in Annex I are not included.

5.1 III.1 Does the application use biometric data?

Following Annex III sub 1, a high-risk application of an AI system is biometrics. Which is covered by QIII.1a. See Figure 23.

In case biometric data is used, the user is redirected to Q5.e.1 to check whether the application is potentially prohibited. If no biometric data is used, the user is forwarded to QIII.2.1.

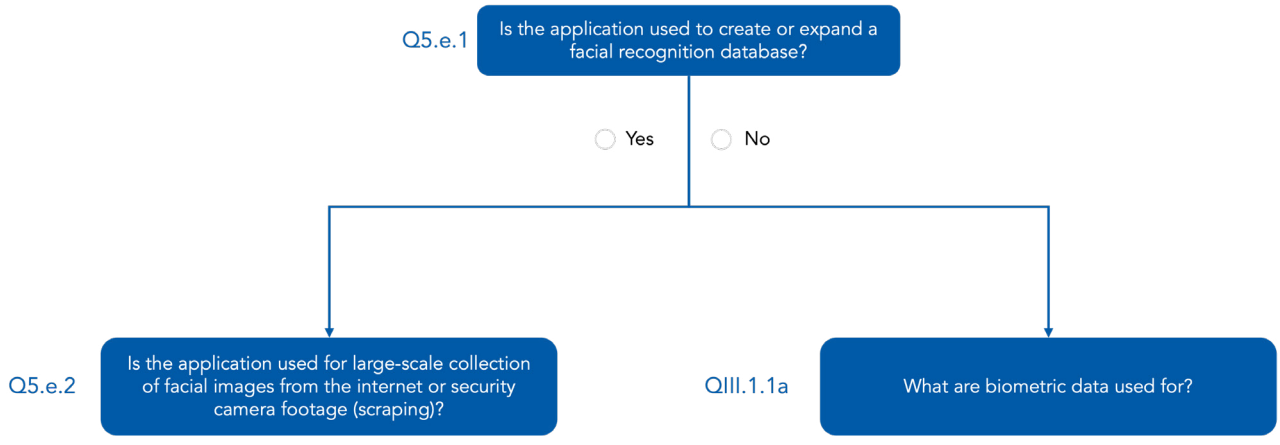

Figure 23 - Q5.e.1 determines whether the application is used to create or expand a facial recognition database.

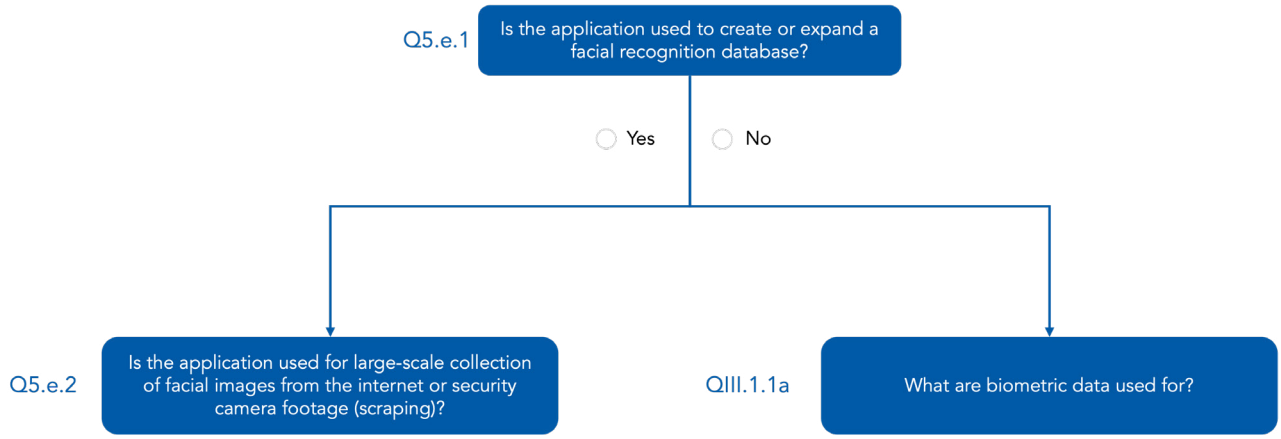

Figure 24 - Q5.e.2 determines whether the application is used for large-scale collection of facial images from the internet or security camera footage (scraping).

To assist users, the following remark is provided:

Biometric data are unique characteristics of a person such as images or video footage of faces, iris scans, fingerprints, voice sounds, handwriting, and other unique characteristics of bodies and/or behavior. Typed text is not biometric data, but someone's unique way of typing (keystrokes) is.

Q5.e.1 Is the application used to create or expand a facial recognition database?

Following article 5 sub 1 under e, an AI system is prohibited if it "create[s] or expand[s] facial recognition databases through the untargeted scraping of facial images from the internet or CCTV footage". Q5.e.1 determines whether the latter is at stake. See Figure 24.

In case a facial recognition database is created or expanded, the user is redirected to Q5.e.2 to check whether the application is potentially prohibited. If no facial recognition database is created or expanded, the user is forwarded to QIII.1.1a.

Q5.e.2 Is the application used for large-scale collection of facial images from the internet or security camera footage (scraping)?

Following article 5 sub 1 under e, an AI system is prohibited if it "create[s] or expand[s] facial recognition databases through the untargeted scraping of facial images from the internet or CCTV footage". Q5.e.1 determines whether the latter is at stake. See Figure 25.

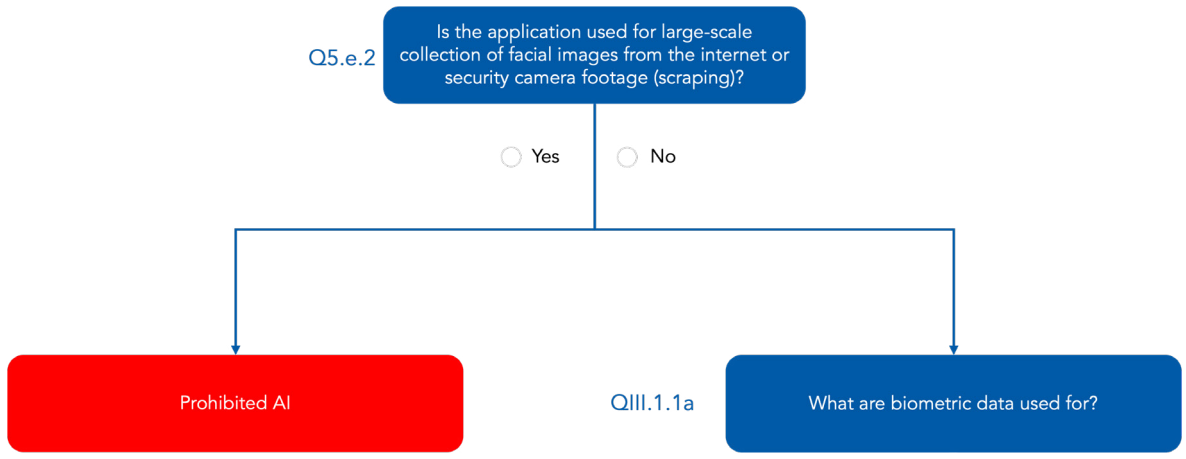

Figure 25 - Q5.e.2 determines whether the application is used for large-scale collection of facial images from the internet or security camera footage (scraping).

Will be updated soon

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