Responsible Use of AI Assistants in the Public and Private Sectors

Netcompany

Table of Contents

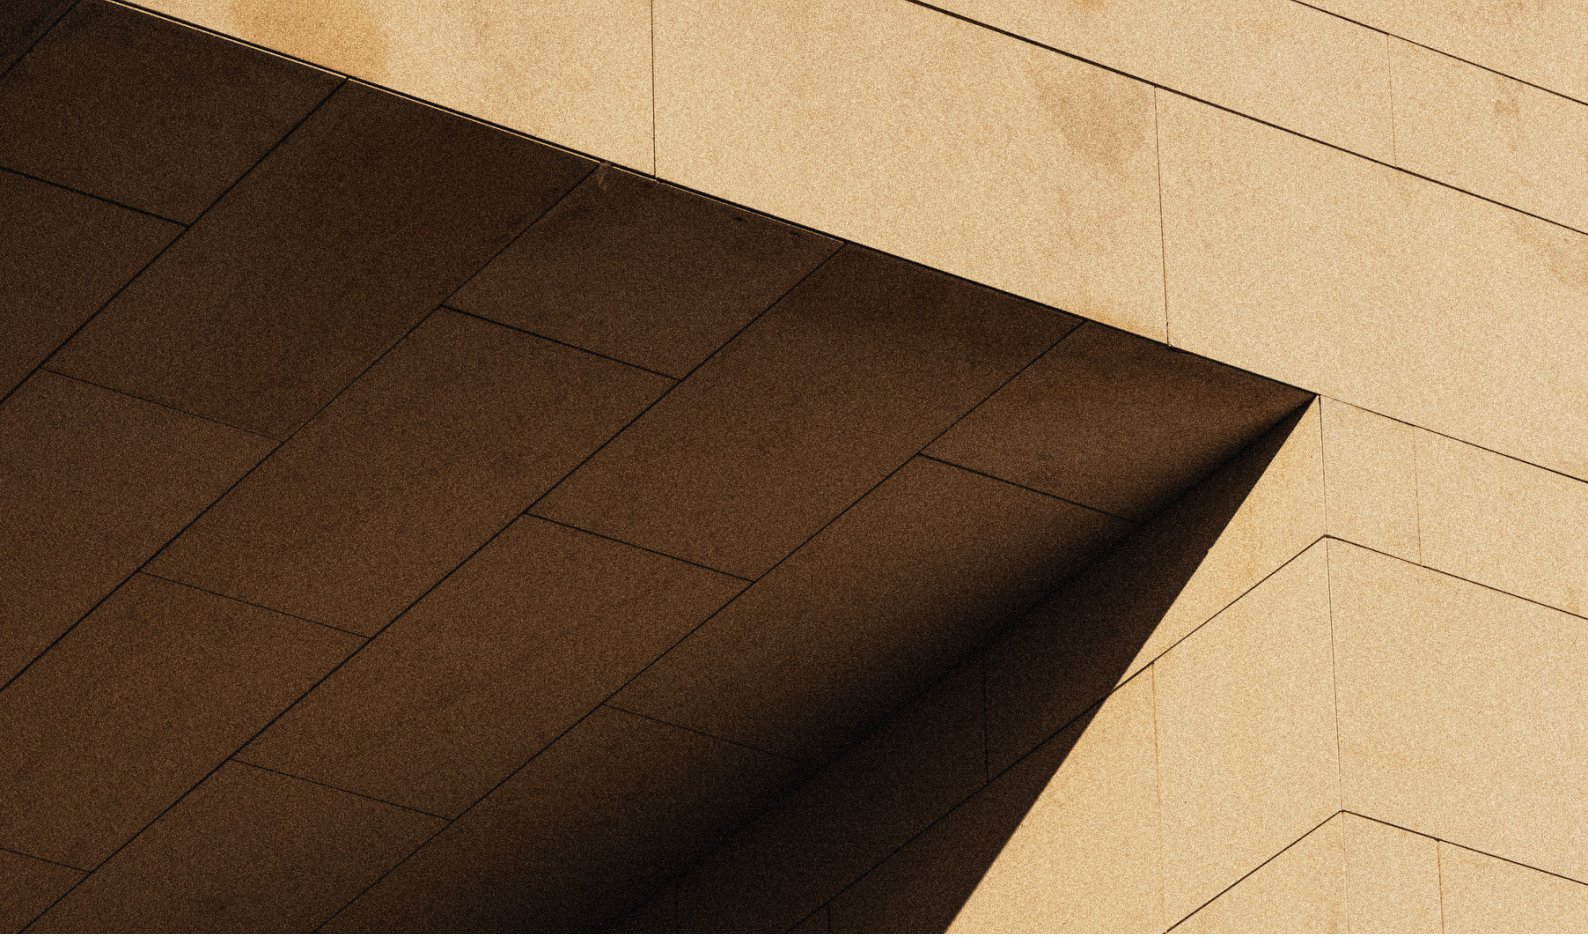

1 Introduction and purpose

In today's digital age, technological advances are increasingly challenging our way of thinking, working, and living. With the rise of artificial intelligence as one of the major drivers of transformation, we are now facing new challenges, as well as unprecedented potential to create better and more effective solutions across different sectors.

In Denmark, we have a unique opportunity to leverage the potential of technology for the benefit of society as a whole – which is at the same time a necessity to ensure our competitive capacity and enhance our welfare. At a time of unrest and upheaval in the world and with rising military and economic tensions, the ability to develop and adopt technologies – in particular, artificial intelligence – will play an increasingly vital role across Europe.

Navigating this new reality – in a legal, technical, and organisational sense – requires a common compass that ensures the development, implementation and use of technology is executed in the best, most innovative and responsible way. It is a journey that we historically have best travelled together.

Fundamental digital infrastructure for Denmark is built in collaboration between public authorities and private-sector suppliers. With the intent of providing guidance and in-depth descriptions of how AI Assistants can be developed, implemented and used, both innovatively and responsibly within the public and private sectors in Denmark, key players from DI Digital, led by Netcompany, as well as Kammeradvokaten – the Legal Adviser to the Danish Government, Microsoft, Dubex and Trifork; together with digitally leading public organisations: the Agency for Digital Government, the Danish Business Authority, KOMBIT and Udbetaling Danmark; and organisations representing the financial sector: ATP, Jyske Bank, PFA, Spar Nord and Topdanmark have entered into a unique collaboration.

These organisations have come together to develop this white paper, which provides a tangible approach to the use of AI Assistants within existing legislation, such as the GDPR and the EU AI Act (referred to in this document as 'The AIA'). It touches on key challenges, such as the risk of disinformation in the form of hallucinations and bias, which require strict quality assurance and robust risk management.

It also provides addenda, which present examples of previous implementations of AI Assistants from the participating organisations, which can provide inspiration and support for future solutions. These examples serve as insight into how technology is successfully used in practice to create optimal value for individuals and businesses.

The white paper marks a significant milestone in the use of AI in Denmark and highlights how cooperation between the country's leading digital stakeholders can lead the development of scalable, reliable and legal use of AI. It is a tool that can help organisations ask the right questions, make informed choices and build on shared experiences.

1.1 Target group

This white paper is designed to be used in both public and private sectors. It is aimed at both senior management, who ensure strategic anchoring and resource allocation, and project teams, responsible for implementation of AI assistants – whether project managers, business representatives, the AI development team, IT architects or legal advisers.

1.2 Terminology

Term	Explanation
AI-assistant	A system or part of a system that uses a language model to support and automate a business process (see section 2).
User	A person who interacts with an AI Assistant. The user may be internal or external to the organisation. Internal users are referred to as 'employees', while external users are referred to as 'citizens'. In the white paper, the term 'citizens' also includes customers of private organisations.
Platform	A technical solution that enables the development and implementation of AI Assistants.
Cloud	Internet-based IT services provided by third parties. Powerful language models require resource-intensive infrastructure and are typically accessed through the cloud.
Domain	An organisation's business area.
Hallucinations	The generation of factually incorrect information by language models, which may be convincing to the user [1].
Logging	Automated recording of key system events that provides insight into the performance, security and patterns of use of an AI assistant. Logs are used for maintenance and problem-solving, as well as documenting the use of the AI assistant and its data processing.
On-premises	Local IT services operated by an organisation or provided by a direct supplier to the organisation. Smaller language models require less infrastructure and can therefore be hosted on-premises.
Prompt Engineering	The process involving the design of instructions that effectively guide language models to provide high quality, relevant and informative results [1].
Language Model	A class of generative model that is trained on a large data set. They typically generate text based on a prompt [1].

1.3 Scope

The white paper focuses on AI Assistants based on Generative AI, primarily language models, and covers the use of other AI types such as predictive models. It assumes the organisation has already identified AI Assistants that create business value and can be supported by language models.

The white paper addresses relevant obligations under the GDPR and AIA as these regulatory environments are relevant to all organisations using an AI assistant. It is not an exhaustive list of all legal obligations; for example, AI assistants that the AIA classifies as being of unacceptable risk are not covered by the white paper.

Different organisations have their own ethical guidelines and standards. It is outside the scope of the white paper to define general ethical guidelines.

2 What is an AI Assistant?

AI Assistants are systems or parts of systems that use language models to carry out parts of a business process with a certain degree of autonomy. They generate responses based on the information and prompts they receive, which can facilitate interaction with other systems, optimise decision-making and create more flexible, data-driven solutions.

An AI Assistant can thus handle tasks that normally require human assessment by using tools and data from existing business systems. This could include drafting minutes, updating information based on email correspondence, providing support for legal advice or offering chatbots for personalised customer service. They are therefore typically implemented with the aim of performing specific tasks more uniformly, efficiently and at higher quality.

AI Assistants can be developed with standalone tasks in mind or as part of complex, multi-assistant working procedures, provided they occur within the framework of applicable legislation and meet specific transparency requirements.

Below are a set of principles that the AI assistant must follow regardless of the business domain. The principles constitute an informative basis that organisations can use to define requirements for their AI assistants.

Principle	The AI Assistant presents the information basis for their responses in an understandable way.
Rationale	When the user can see the information on which the model bases its responses, it is easier to assess the quality of the results and identify any hallucinations. A typical approach is that the model describes its information basis or gives references to the source.
Principle	The AI assistant only uses data that is considered relevant to a specific prompt.
Rationale	Using relevant data ensures that the AI assistant's responses remain focused and accurate. This reduces the risk of introducing irrelevant or misleading information and thus improves the quality and reliability of the responses. In addition, it optimises efficiency by limiting the quantity of data to be processed.
Principle	The AI assistant performs essential tasks independently and informs the user if a task is beyond its responsibilities or it requires assistance.
Rationale	Operating independently maximises the AI assistant's utility and the user's trust in the system. When the user is aware of limitations or insufficient information, misunderstandings are minimised, and when the information is retrieved from both a prompt and a knowledge base, potentially erroneous responses are avoided.
Principle	The AI assistant understands the context of a prompt, as well as previous, relevant messages from the same conversation.
Rationale	By integrating previous interactions, the AI assistant can provide more nuanced and relevant responses. This reduces the need for repetition and makes communication more natural and effective.

3 Nine points of the process

This section presents a systematic and iterative process that guides organisations in the development, implementation and maintenance of secure, effective and responsible AI assistants. The process comprises a total of nine points that are important activities when working with safe and responsible AI. These include everything from initial planning to ongoing operation and optimisation (Figure 3.1). Although the process is similar to that typically seen when implementing IT systems, this process places particular emphasis on those aspects that are of crucial significance to AI assistants.

The relevance of these points inevitably depends on the needs and regulation of each organisation and the respective sectors. However, it is essential that organisations consider and document the reasons why they opt out of specific points.

The points are organised in a logical sequence but should be considered as interconnected and mutually influencing. The approach is iterative, as organisations will often need to repeat and reassess as understanding grows, new challenges arise or the scope of application for the AI assistant expands. This could, for example, be changes to the legal framework, while new data may result in further considerations of legislation and sensitivity. Organisations must thus stay updated on legal matters to ensure they comply with the principle of lawfulness.

Figure 3.1: The nine points for AI implementation

The plan comprises of the following nine points:

1. Define the AI assistant's use case
2. Establish the AI assistant on a flexible technical platform
3. Assess necessary data and data processing
4. Address the legal framework
5. Set boundaries for the abilities and responsibilities of the AI assistant
6. Build structured quality assurance
7. Measure and store relevant data on the use of the AI assistant
8. Plan organisational implementation and training
9. Establish follow-up and support structures

Each individual point contributes to different phases of the AI assistant's lifecycle, from planning and design to implementation and operation. By following these points for iterative implementation and maintenance, organisations can adopt a holistic approach that takes into account technical, organisational, legal and business aspects. This flexible approach enables ongoing learning, adaptation and optimisation in line with the development of technology, organisation, legislation and the outside world. Each point is expanded upon in the following sections.

3.1 Point 1 – Define the AI assistant's use case

Valuable use cases for AI assistants can include full or parts of work processes that require collaboration between multiple people within the organisation and which cannot be solved with simple functions. While AI assistants can perform standardised tasks, the generative nature of language models allows them to handle more dynamic processes in which decision-making is required along the way. It is important to consider in what way your organisation's current technology might be obsolete. Language models offer innovative solutions to tasks not covered by existing software solutions, while also harnessing the potential of AI assistants.

Implementing an AI assistant expands the organisation's digital capabilities. To ensure successful integration, the use case must transform the AI assistant's vision into a specific description that forms the basis for further activities. This use case should be rooted in the specific issue that the AI assistant must address, with a focus on the business value it is meant to create. When developing the use case, consideration should be given to four perspectives: process, users, information and tools.

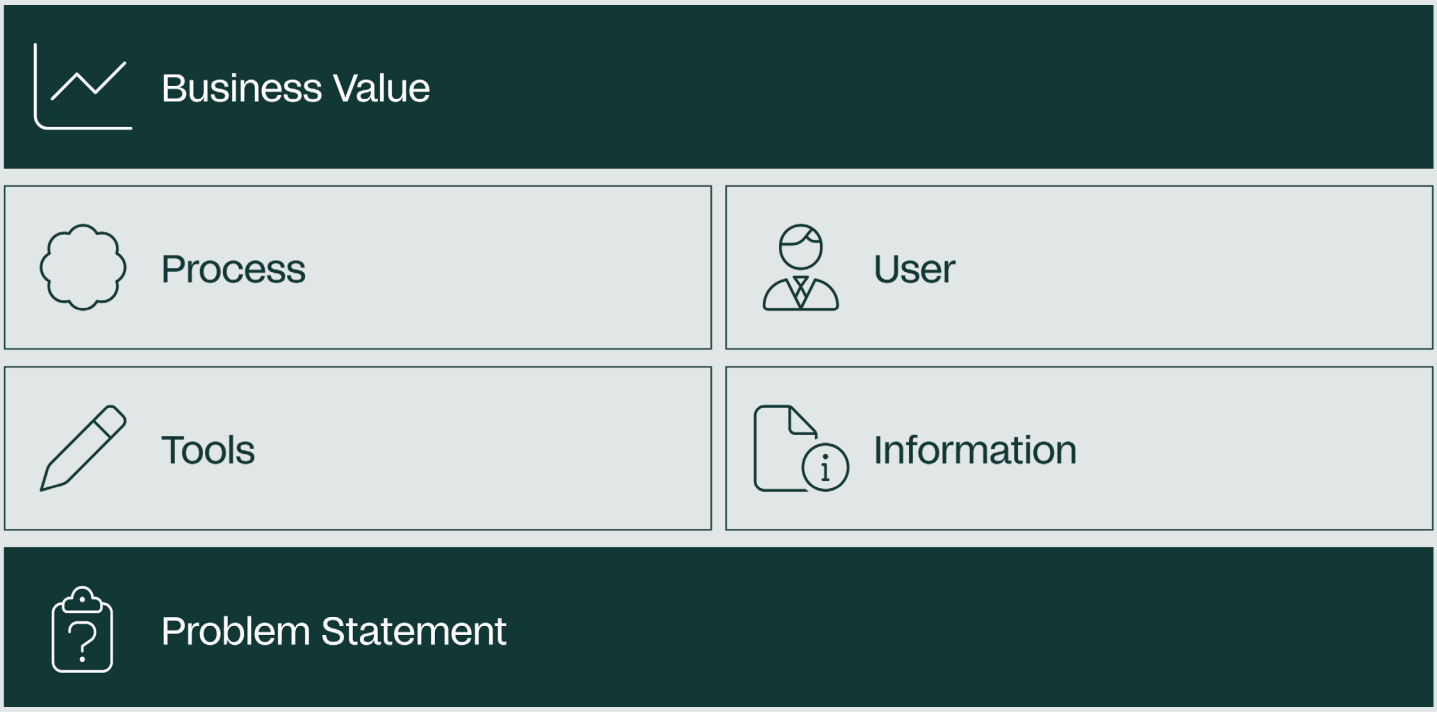

Figure 3.2: The four perspectives – a use case described from issue to business value

Problem statement
A use case should have a clear objective that the organisation wants to address, such as improving quality, optimising services or improving the efficiency of processes. The problem does not necessarily need to be challenging; it could also be an opportunity, for example, to upgrade an older digital solution or to improve existing processes.

Process
The working process supported by the AI assistant may be simple, such as finding information within a document, or complex, such as creating the analytical framework for an assessment. The current human working process is mapped out, after which the AI-supported process is described.

Users
It is important to consider who the users are, how they access and interact with the AI assistant and how the results impact their work. Interaction with the AI assistant may vary depending on whether the users are internal employees who can be trained in its use or individuals who are simply guided by it, while implementation can be hindered by resistance to new tools. Section 3.8 concerns the training of users.

Tools
The functions used by the AI assistant to perform its tasks may involve integration with other systems for calculations or data modifications. An AI assistant that performs assessments may, for example, need access to a model for quantitative analyses because language models are not suitable for this.

Information
The AI assistant must have relevant knowledge available to be able to perform its tasks correctly. This may include data from other systems, guidelines, etc. With correct information, the model is better equipped to operate within its domain, reducing the risk of hallucinations [2]. Further considerations on data and data processing are reviewed in section 3.3.

Business value
Defined KPIs, such as response times, error rates, productivity boosts and customer satisfaction, for the AI assistant's purpose are critical to demonstrating its value. By establishing a pre-implementation baseline for performance, its effect can be better assessed. While it may be challenging to set precise goals for new technologies, subjective factors such as the level of efficiency and satisfaction experienced should also be measured. In addition, any costs for operation and support must be taken into consideration.

The above forms the foundation for the subsequent implementation process. It provides a clear direction and understanding as to how the AI assistant integrates and creates value in a specific context, while addressing potential challenges and risks proactively.

3.2 Point 2 – Establish the AI assistant on a flexible technical platform

To promote security and efficiency, AI assistants should be centralised in an AI platform with a robust security model. This prevents unauthorised solutions and allows reuse and integration with other systems when multiple use cases need to be covered.

An organisation should strategically consider how it wishes to work with AI assistants. There are basically three approaches:

1. **In-house development:** Provides full control and customisation but requires significant resources and technical expertise.
2. **Supplier collaboration:** Balances control and external technical expertise.
3. **Outsourcing:** Rapid deployment but less control over the solution.

An organisation can employ various strategies for using AI assistants. Non-business-critical AI assistants can be purchased as standard products, while core business processes may require in-house development, and when that AI assistant is incorporated into an organisation, it becomes part of the organisation's overall system landscape.

A platform offers multiple benefits, such as lower costs, shorter time-to-market and increased innovation through component reuse, while the choice of a standard platform or tailored solution depends on the needs of your organisation. A standard platform enables faster implementation plus lower initial costs and support, while a tailored platform offers full control and customisation to specific needs, but with longer development times and higher costs.

By considering the long-term strategy and goals of the organisation, simple AI assistants can be developed on a standard platform, while domain-specific AI assistants can benefit from a customised platform. The approach chosen affects your responsibilities in relation to AIA, see section 3.4.1.

It should be possible for a platform to integrate new language models so that the organisation can switch models in line with technological development. Figure 3.3 illustrates the components in a system landscape where the AI platform is loosely connected with the model integrator, model providers, knowledge base and potential support features.

By thinking one's systems into this landscape, enables the possibility to replace language models in line with the development of technology and changes of prices, as well as the reuse of AI assistants across mobile apps, websites, desktop applications, etc. A review of the components of the system landscape is presented below.

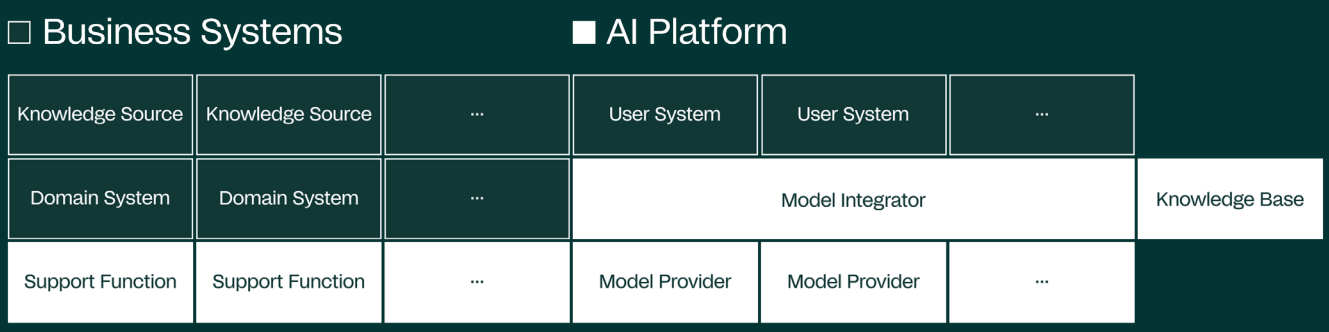

Figure 3.3: The components in a system landscape centred on the model integrator, which exposes AI assistants.

Model provider
Model providers are platforms that make language models available to AI assistants through system interfaces. These providers may be cloud-based provided by an external supplier or on-premises solutions provided by the organisation itself. When selecting a model provider, the organisation needs to consider whether it can handle the number of anticipated prompts that the model will receive, as well as ensuring that the model provider handles data correctly and complies with the GDPR. When switching models, the new model provider should be quality-assured in the same way as the previous provider, see section 3.6. The model provider does not necessarily have to be the provider of the AI assistant as laid down in the AIA, see section 3.4.1.

Model integrator
The model integrator orchestrates the interaction between various components, such as user system prompts, selection of the appropriate model from a model provider and use of data from the knowledge base. The AI assistants are exposed by the model integrator as services and are not directly connected to the model providers. This ensures that the AI assistants are independent of each model provider while also having access to relevant business systems.

Knowledge source
Knowledge sources provide domain-specific information to the AI assistant that it uses to perform its tasks. The information can be both structured and unstructured, and some information requires preparation such as embedding.[1] When integrating knowledge sources, it must be ensured that data is accurate, up to date and relevant (see section 3.3) and that data is available and in a format that can be processed by the model integrator.

Domain system
Domain systems are IT systems within an organisation that support business processes. An AI assistant uses a domain system to perform actions such as creating, modifying or removing information. The action that an AI assistant performs in the system must be displayed through a system interface. The AI assistant must not perform any action on behalf of the user that the user is not entitled to perform. For example, if the AI assistant attempts to delete information to which the user has read-only rights, the action must be rejected.

User system
User systems are the applications through which users interact with AI assistants. These may be chat-based applications or existing systems in the organisation. A user must not have access to an AI assistant that performs tasks which that user does not have access to perform. The user system can present the AI assistant as a chat interaction or use the AI assistant in the background to handle a process automatically. To support good organisational implementation, it is important for the interaction with the AI assistant to be user-friendly. Similarly, consideration should be given to how critical tasks can continue to be performed if the AI assistant is unavailable. If the AI assistant makes decisions or handles the processing of individuals, human oversight must be taken into account, see section 3.4.1. In these cases, a user system could be used.

Knowledge base
A knowledge base is a centralised repository for storing and handling information that the language model uses to answer prompts. Knowledge bases are typically implemented through a database. The knowledge base typically contains data from knowledge sources that needs to be prepared before use. When selecting a knowledge base, it is important to consider how the data is connected and how the model integrator can quickly and efficiently find the necessary information.

Support system
Support systems include a variety of functions and technologies that ensure AI assistants perform optimally and in accordance with organisation's requirements and standards. They are typically used to build, train and fine-tune models.

A system can have multiple roles within the system landscape. For example, a case processing system might be a knowledge source from which an AI assistant retrieves information about a specific case, a domain system in which the AI assistant updates information on the case and also the user system itself, where the case officer interacts with the AI assistant. When a system has multiple roles, it is important to consider the different requirements this may create.

[1] Embedding: A method where text is transformed into vectors that represent the context of the text and can be used to perform subsequent searches by using the embedding of a prompt.

3.3 Point 3 – Assess necessary data and data processing

Certain tasks require the AI assistant to have task-specific knowledge such as guidance, case information or similar data that varies in format, sensitivity and structure. To ensure trust and to avoid data leakage, data management must be structured, secure and documented. Data must be continuously updated to ensure accurate and relevant background information for the AI assistant, and processes for management and administration are required by the AIA to ensure data quality (see section 3.4.1).

Generally, there are two types of data: structured and unstructured data.

Structurered data: Information in a predefined format, such as customer information, transaction data and product data. Structured data can be deployed in tables where the values are organised in columns.

Unstructurered data: Information whose content cannot be categorised or indexed in a consistent manner. This is typically text documents, websites, videos and images. Language models are primarily trained on unstructured data and are therefore effective at handling natural language.

Structured data is often easier to handle systematically than unstructured data, and the degree of data structure can vary. For example, a collection of documents may have uniform headings that create some structure, while structured data may contain fields with images or formatted text. To prepare data, the following elements are proposed as relevant to ensure quality:

Correctness: How accurate and reliable the information is in relation to the domain in which it will be used. A manual review is frequently required to ensure correctness of unstructured data because the underlying context and interpretation are important. Structured data is easier to validate due to the predefined format and often can be automatically checked for errors and inconsistencies.

Currency: How relevant and up to date the information is. Unstructured data may contain time-sensitive information that requires regular updating, for example documentation of a system under development. Structured data may be more easily integrated with automatic update systems. For both data types, it is important to assess how often data is updated or deleted.

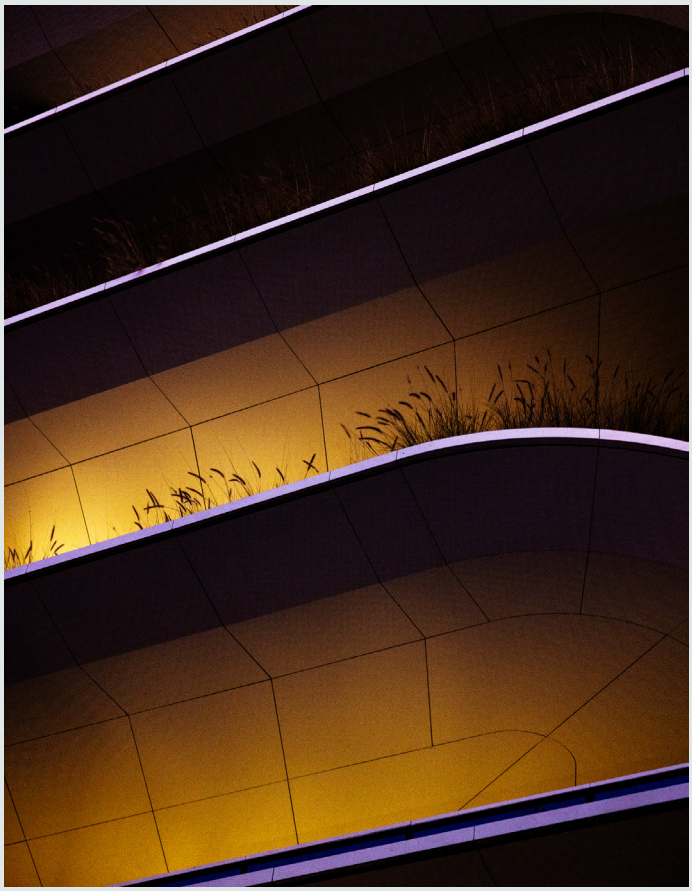

Context: How clear and explanatory the information is regardless of its context. Unstructured data is typically rich in context, while structured data often requires supporting documentation or metadata for proper interpretation.

Confidentiality: How sensitive the information is, such as personal data or business-critical strategic data. It can be a challenge to automatically identify confidential information in unstructured data unless the confidential information is known and simple. Structured data with confidential fields is often simpler to prepare.

Availability: How easily and simply data can be accessed in relation to security and uptime. Data availability can be challenging for AI assistants because information is often spread over different systems, with different access controls and different expectations as to when the system can be accessed.

Bias: How representative and accurate data is. Unstructured data may have hidden biases that can be difficult to spot out of context. Structured data may also be biased, especially if data collection is limited. Bias can often be measured and corrected in data, but it requires attention to collection methods and representativeness.

With reference to the above focus on data quality, the AI assistant's data work should focus on data security as well as existing practices such as the CIA triad.[2]

[2] The CIA triad consists of three key information security elements: confidentiality, integrity and availability. Confidentiality: Data must only be accessible to authorised persons and be protected from unauthorised access. Integrity: Data must be accurate, complete and protected against unauthorised modification. Availability: Data must be available and usable whenever authorised users require it [9].

3.3.1 Data processing

When an AI assistant prepares and sends data to a tool or underlying language model, it is important to implement robust security measures and to comply with data protection rules such as the GDPR, see section 3.4.2. In the case of data processing, a distinction can be drawn between public data and personal data.

Public data: Information freely available to the public, such as public records, open data sets from governments and company records or information. Public data does not require special protection or consent to be processed.

Personal data: Personal data is categorised by sensitivity under personal data legislation. A distinction is made here between ordinary personal data, such as name, address and telephone number, and sensitive personal data, which is characterised by information such as race, religion, health and sexual orientation. The GDPR further regulates children's data and criminal data, while the Danish civil registration number (CPR) is separately regulated under the Danish Data Protection Act [3].

An organisation may choose to classify certain data, such as competition-critical process descriptions, as confidential, even if it does not contain personal data. In such cases, a data breach may have a competitive impact, which makes it important for organisations to have systematic protection measures in place for such data. When working with data, it is important to think in terms of the principles of data protection by design and data protection by default, see [4].

When considering whether a model should be cloud-based or on-premises, it is important to think about how the data will be processed. Public data may be shared through a cloud-based model, but if personal data is sent to a cloud solution, the data processing by the provider must be accurate and documented. On-premises language models reduce dependency on third parties and allow businesses to implement customised security policies. These language models are typically smaller, which limits the complexity of the tasks that they can perform. There is therefore a trade-off between confidentiality and quality when choosing between a cloud or on-premises model.

The AI assistant should only have access to data that is relevant to the task and to which the AI assistant's user has access (see section 3.5). There are several methods that can be used to minimise the data processing risk when the model is provided via cloud.

Retrieval-augmented generation (RAG): A method in which relevant information is prepared in advance and is received from a knowledge base by request. The information is deployed as context in the prompt. When using RAG, the language model generates more accurate and context-based responses. Unlike traditional training of the language model, where data is integrated into the model, data exists independently from the model when using the RAG method. Data can therefore be more easily updated without having to retrain the model. RAG simultaneously lowers the risk of hallucinations because information becomes part of the context of the model [2].

Systematic data masking: Known, sensitive data in a prompt is replaced with temporary values before it is processed. When the processing is completed, the original data is reinserted into the response. Data that is not known in advance cannot be systematically identified. Efficient data masking, where information cannot be traced back to individuals, avoids the processing of personal data.

AI-based anonymisation: Use of the on-premises model to identify sensitive data, which is then masked. Both language models and specialised masking models can be used here. There is a risk that some personal data may not be properly identified if information is structured in a way that the model has not seen before, leading to it inadvertently being shared with the cloud provider.

Data minimisation: A method by which information related to an individual or their situation is sent without the personal data being forwarded. This may be the age group, legislation related to the individual's situation or terms and conditions of a service. The AI assistant receives relevant information but no individual personal data. This enables an AI assistant to support sensitive processes, such as case management, without sharing personal data with a cloud-based model. Data minimisation means that context is removed from a prompt, which can cause the model's responses to be more abstract.

An organisation has a duty to document its approach to data processing. By implementing robust security measures, complying with data protection rules, managing risks and implementing management and administrative processes around the AI assistant's data, the foundation is created for reliable integration into the organisation's processes and thus ensuring valuable support without compromising data security.

3.4 Point 4 – Address the legal framework

Based on point 1, see section 3.1, and the business and technical clarifications that have been made, an organisation must first identify the relevant legal requirements that the AI assistant must comply with prior to its use. As mentioned in point 1, this step should be performed continuously throughout the lifecycle of the AI assistant, particularly in the event of changes in use.

In particular, there are two set of rules that must always be considered and taken into account – the General Data Protection Regulation (GDPR) [5] and the Artificial Intelligence Act (AIA) [6], which are addressed in this section. A checklist has been prepared, included as Appendix A, for the purpose of clarifying and operationalising relevant requirements of the GDPR and AIA.

However, other relevant legislation should also be considered depending on the specific use of the AI assistant.

Health legislation, consumer protection legislation, rules on confidentiality, financial legislation, intellectual property law, management rules, procurement rules, NIS2, etc. may be applicable here. It will also include any legislation that the AI assistant has to answer questions about, which must also be identified and handled as part of the mapping of the legal framework for the assistant.

The checklist in Appendix A can be used in both in-house development cases and when the AI assistant is purchased as a standard solution. The checklist can also be used if you wish to ensure that an assistant follows good practice for responsible development and use of AI, regardless of whether or not the assistant is covered by the AIA.

3.4.1 The AIA

The AIA applies to 'AI systems'[3] in accordance with the Regulation. It is presumed that the AI assistant is considered an AI system in the sense of the AIA, and it is also assumed that the AI assistant falls within the scope of the AIA. In general, there are two factors that need to be clarified in relation to the AIA on this basis:

1. Identify the AI assistant's risk level; see section 3.4.1.1
2. Identify one's own role under the AIA; see section 3.4.1.2

These circumstances are elaborated on below. Please also refer to Appendix A, under 1.2.1 and 1.2.2, in relation to the requirements for providers and deployers of high-risk AI assistants, respectively.

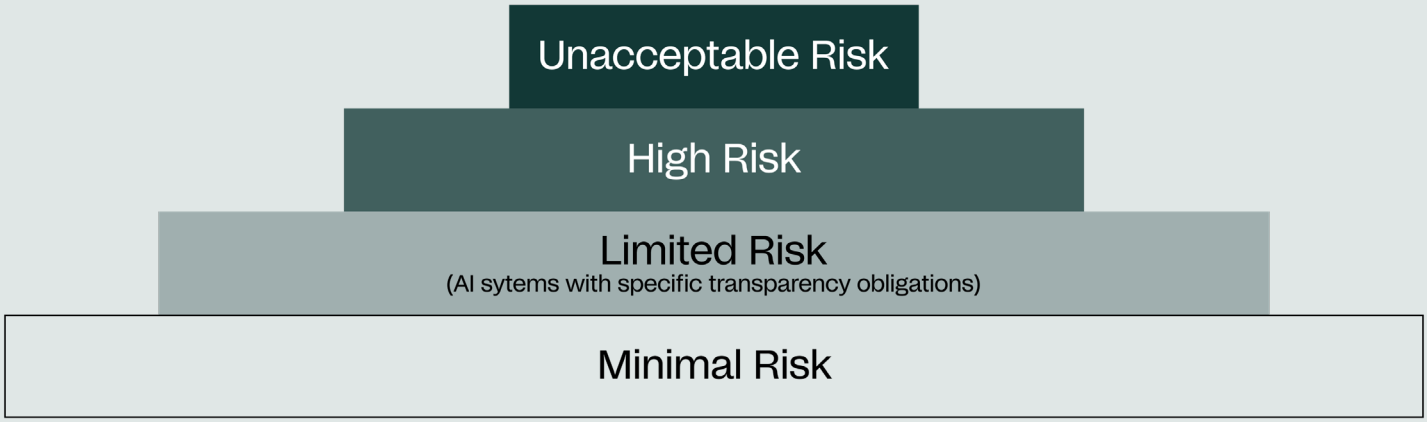

Figure 3.4: The four risk levels defined in the AIA [7].

3.4.1.1 Identify the AI assistant's risk level

The AIA operates with a risk-based approach, and there are therefore different obligations under the AIA depending on the risk inherent in the AI system and the role of the organisation. The risk-based approach in the AIA is illustrated in Figure 3.4.

Prohibited AI-systems
Several types of AI systems are prohibited under the AIA because they are considered to pose a significant threat to human safety and the rights or dignity of citizens. Examples include social scoring systems and voice assistant technologies that promote dangerous behaviour. It must therefore be investigated and clarified whether the AI assistant is used in a manner prohibited by the AIA, such as the assistant using manipulative or misleading techniques to persuade a user to make a particular decision, to the detriment of the user.

High risk
An AI system is considered high risk if the following conditions are met:

1. The AI system is intended to be used as a safety component in a product or is in itself a product subject to EU harmonisation legislation, as laid out in Annex I to the AIA.

2. The safety component or product will be subject to a conformity assessment performed by a third party for the purpose of the placing on the market or putting into service of the product in accordance with EU harmonised legislation, as laid down in Annex I to the AIA.

[3] Article 3, point (1), of the AI Act: 'AI system' means 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' [6].

In addition, the AI systems referred to in Annex III to the AIA are also considered high risk. An AI assistant can thus fall under the high-risk category if it is used, for example, in:

- critical infrastructure (e.g. transport) that could endanger citizens' life and health.
- education or vocational training that may be crucial for an individual's access to education and vocational courses (e.g. examination assessment).
- employment, workers' management and access to self-employment (e.g. CV sorting software for employment procedures).
- essential private and public services (e.g. credit rating that denies citizens the opportunity to obtain a loan).
- law enforcement, which can interfere with fundamental human rights (e.g. evaluation of the reliability of documentation).

If an AI assistant is considered a high-risk AI system, it is subject to strict obligations before it can be placed on the market. The requirements are as follows:

- Appropriate risk assessment and risk mitigation systems.
- High quality of data sets that provide data for the system to minimise risks and discriminatory results.
- Logging of activity to ensure traceability of results.
- Detailed documentation of all the necessary information about the system and its purpose so that the authorities can assess compliance with the requirements.
- Clear and sufficient information for the deployer.
- Appropriate measures for human oversight to minimise risk.
- High degree of robustness, security and accuracy.

Please refer to the checklist in Appendix A, under 1.2.1 and 1.2.2, for more information on requirements related to high-risk assistants.

Limited risk

Limited risk refers to the risks associated with lack of transparency in the use of artificial intelligence. The AIA includes specific transparency obligations to ensure people are informed when necessary, thereby promoting confidence. For example, when using AI systems such as chatbots, users should be made aware that they are interacting with a machine so that they can make an informed decision to continue or step back. Providers must also ensure that AI-generated content can be identified. In addition, AI-generated text published with the aim of informing the public about issues of public interest must be labelled as artificially generated. This also applies to audio and video content that constitutes deep fakes.[4]

Assistants that fall under the restricted risk category must therefore comply with Article 50 of the AIA.

Minimal or no risk

The AIA allows free use of artificial intelligence with minimal risk. This includes applications such as AI-enabled video games or spam filters. The vast majority of AI systems currently in use in the EU fall into this category.

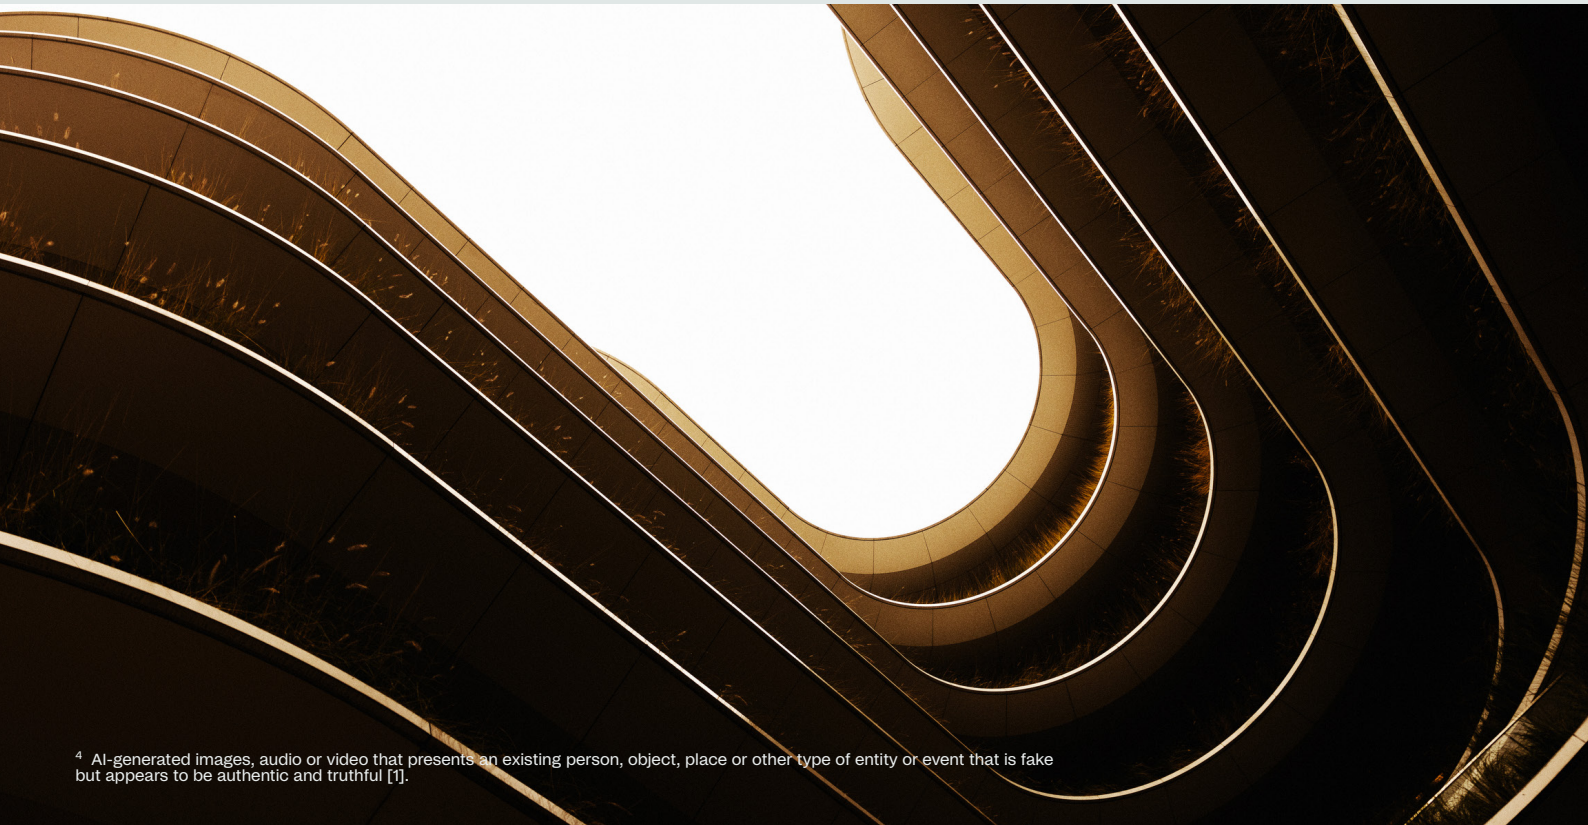

[4] AI-generated images, audio or video that presents an existing person, object, place or other type of entity or event that is fake but appears to be authentic and truthful [1].

3.4.1.2 Identifying your role under the AIA

The scope and type of your obligations under the AIA depend on your role. There are several roles under the Act, but typically the question will be whether your organisation is a provider or deployer of the AI assistant:[5]

Provider
A natural or legal person, public authority, agency or other body that develops an AI system or a general-purpose AI model or that has an AI system or a general-purpose AI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.

Deployer
A natural or legal person, public authority, agency or other body using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity.

The vast majority of the obligations under the AIA rest with providers of high-risk AI systems (Appendix A, under 1.2.1). Please be aware that a number of obligations are the responsibility of other parts of the AI systems' value chain, such as providers of:

General-purpose AI-model
An AI model, including a model where such an AI model is trained with a large amount of data using self-supervision at scale, that demonstrates significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development or prototyping activities before they are placed on the market.

General-purpose AI-system
An AI system which is based on a general-purpose AI model and which has the capability to serve a variety of purposes, both for direct use as well as for integration in other AI systems.

Downstream provider
A provider of an AI system, including a general-purpose AI system, which integrates an AI model, regardless of whether the AI model is provided by themselves and vertically integrated or provided by another entity based on contractual relations.

When 'fine tuning'[6] a language model that supports an AI assistant, it is appropriate to consider whether you are a provider of a general-purpose AI system or if you are a provider of a high-risk AI system. You must therefore be aware of the obligations in the AI systems' value chain, as described in Article 25 of the AIA, and assess under which part of the value chain an assistant falls.

[5] Definitions specified in this section are set out in Article 3 of the AIA [6].
[6] Further training of a model on top of a specific data set with the aim of specialising the model.

3.4.2 GDPR

The GDPR applies if personal data is processed when developing and/or using the AI assistant. If your organisation uses an AI assistant that you have either developed or purchased as a standard service (Software as a Service solution), you will usually be the data controller for the processing of the personal data in the AI assistant. You must therefore be aware of the obligations of data controllers in the GDPR.

Below are the typical requirements and themes that are in play when you, as data controller, process personal data when developing/using an assistant.

Get an overview of the processing of personal data
You should start by getting an overview of the data processing, including what personal data is processed about which individuals for what purposes throughout the solution's lifecycle, including during the development, testing and operation/retraining of the solution. This applies to both input data, including any personal data in prompts, as well as output data. Be aware of whether use of the assistant involves generating new personal data in its output data; for example, the assistant might deduce or infer sensitive personal data about the user or other people through profiling.

Manage the data protection rules throughout the lifecycle of the solution
A fundamental requirement is that the data controller must be able to document compliance with the data protection regulations. As a result of the requirement for data protection by design and by default in Article 25 of the GDPR, the data controller must factor in the data protection rules and manage risk throughout the lifecycle of the AI assistant, i.e. right from the design phase until monitoring of the solution in operation. This applies regardless of whether these are standard solutions ('off the shelf') or solutions specially developed for a particular task by the data controller.

When processing personal data during the development and operation of the AI assistant, the basic principles of, among others, lawfulness, fairness and transparency, purpose limitation, data minimisation and data quality in Article 5 of the GDPR must be observed. You must also ensure that a basis for processing (legal basis) exists for the processing of personal data. Companies that act as data controller will often have to consider the use of consent from data subjects, contracts or application of the rule of legitimate interest to process non-sensitive personal data under Article 6 of the GDPR. Public authorities will often find the legal basis in Article 6(1)(e) of the GDPR on the processing of personal data as part of the exercise of official authority or a task in the public interest, as well as special legislation in the area in which the assistant is to be used. The more intrusive the processing is for the data subjects – i.e. the more the processing has an impact on citizens' economic, educational, social, health or similar conditions – the stricter the requirements for the clarity of the legal basis. If sensitive personal data is processed, this must be pursuant to Article 9(2) of the GDPR, see Article 6 of the GDPR.

Inform the data subjects about the processing of personal data and manage the rights of data subjects
As a starting point, data subjects must be informed of the processing of personal data that the AI assistant involves; see Articles 13–14 of the GDPR. This may occur through a privacy policy or otherwise. Please note that the development and operation of the AI assistant must be regarded as two separate objectives, which means that the data subjects must therefore be notified of both objectives.

The data subjects must also be informed of the existence of profiling and its consequences. And if automated individual decision-making is carried out by the AI assistant under Article 22 of the GDPR, the data subjects must be informed of this

and receive at a minimum meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subjects.

You must also be able to manage the rights of data subjects, including by providing insight into the processing of personal data and being able to comply with the right to rectification, the right to object, the right to erasure, etc.

Be especially careful if you use the AI assistant for automated individual decision-making

A prohibition on the use of automated individual decision-making applies generally; see Article 22 of the GDPR. However, such decisions may be lawful under the provision if the exceptions to the prohibition are met, including if explicit consent is obtained, the processing is necessary for entering into or performance of a contract between the data subject and the data controller, or if the processing is pursuant to EU or national law. In certain situations, however, the right of data subjects to be able to obtain human intervention in the decision etc. applies.

Manage security and handle risks effectively

The processing of personal data when developing and using the AI assistant must meet the fundamental requirements of processing security. Appropriate organisational and technical security measures must therefore be established throughout the lifecycle of the solution and this must be based on a specific risk assessment concerning the rights and freedoms of data subjects; see Article 32 of the GDPR. When preparing the risk assessment, the starting point must be the purpose of using the AI assistant and its limitations, as well as the nature of the personal data involved. For example, using the AI assistant as decision support for decision-making or profiling of citizens, customers or employees is riskier than, for example, providing employees with the opportunity to use the AI assistant to obtain answers to questions on internal rules on the intranet or to use the AI assistant to produce content for marketing material.

If the processing of personal data is likely to involve significant risk to the rights and freedoms of data subjects, you must prepare a data protection impact assessment (DPIA) prior to the processing of personal data in accordance with Article 35 of the GDPR. Such a DPIA is particularly relevant when processing personal data with the use of new technology, and this will normally be the case when developing and using AI solutions. This particularly applies in the processing of sensitive personal data to a large extent, the processing of vulnerable data subjects' personal data or the use of profiling or automated individual decision-making.

The following are among the typical risks associated with developing and using AI assistants that you will need to manage:

- Factually incorrect responses and hallucinations, i.e. the AI assistant producing incorrect or fictional responses (output).
- Arbitrary discrimination (bias), i.e. the AI assistant causing illegal discrimination due to one or more sensitive criteria such as gender or ethnicity.
- Automation bias, i.e. the people who use the AI assistant uncritically follow or trust the AI assistant's output or use it to make a decision because they trust the solution.
- Insufficient processing security, i.e. unauthorised persons gain access to personal data or manipulate the AI assistant to make it work improperly.

It is therefore also important that the necessary instructions and information are drawn up and provided to employees on how to use the AI assistant correctly and safely so that they are aware of the purpose and limitations of the AI assistant, its legal use and the risks involved. It may also be appropriate to perform quality controls on the AI assistant's outputs etc.; likewise, consideration can be given to requiring employees to clearly mark when output has been generated by the AI assistant.

Transfers to third countries and sharing of personal data with the supplier

Assistants that are provided as Software as a Service solutions in a cloud environment may involve transfers of personal data to unsecure third countries, including to subcontractors. You will therefore need to establish whether such transfers of personal data take place to unsecure third countries; if this is the case, you must assess whether this can occur legally within the data protection rules. In practice, this means that a transfer impact assessment (TIA) must be prepared that maps the transfers and assesses their lawfulness.

It will also be relevant to investigate whether you are forwarding personal data to the supplier of the solution for the supplier's own purposes when using the solution, e.g. training the solution, product development, marketing, etc. In such cases, you must assess whether there is any basis for this disclosure and whether the supplier's use of the personal data for its own purposes is consistent with the original purpose of processing the personal data; see Article 5(1)(b) of the GDPR [5].

3.5 Point 5 – Set boundaries for the abilities and responsibilities of the AI assistant

Boundaries must be set in order to limit AI assistants to specific, well-defined tasks, areas of responsibilities and tools in order to minimise the risk of errors and inappropriate use. This is especially important when several specialised AI assistants are intended to collaborate and tasks are automatically distributed between them. Tools can expand the capabilities of the AI assistant from purely linguistic transformation to more complex tasks. While the underlying language models handle natural language, an AI assistant can perform quantitative analyses, for example, if combined with the right tools or models.

Well-defined tasks also facilitate systematic monitoring and improvement of the AI assistants' performance and reliability. Mature models and AI platforms may already have tools to support to support scoping, and the organisation should explore how they can build on top of existing functionality.

There are many methods of setting boundaries for the capability and responsibilities of AI assistants. Here are some of the most widely used.

Prompt engineering
Precise instructions that target the AI assistant's work within its area of responsibility ensure efficient performance of complex tasks. However, prompt engineering alone cannot be used to set boundaries for the AI assistant, as attacks such as prompt injection could ignore such boundaries. Prompt engineering allows the use of various tools that work together with the AI platform, using the following common methods:

- **Single- og few-shot prompting:** The AI assistant is given one or more examples for the desired response to solving the task. This provides the model with context to work from, which can improve the level of precision for certain types of tasks [8].

- **Prompt chaining:** Breaks down complex tasks into simpler steps that are resolved one at a time. Since the resolution of each step builds on the results of previous steps, the AI assistant builds a deeper understanding of the task. This spurs the AI assistant on to a more well-thought-out, nuanced and accurate result.

- **Agent:** A dynamic form of prompt chaining, in which the AI assistant can take advantage of more specialised AI assistants, functions or other tools, called competences. The AI assistant analyses the prompt and then identifies which sets of competences is needed to solve the task. If a task does not fall under one of these competences, the AI assistant must object.

- **Structured prompting:** The AI assistant must return a response in a specific structured, predefined format. This provides more uniform, machine-readable results that can subsequently be validated and used in other contexts, such as function calling, where the AI assistant's structured response is forwarded to a function or as a prompt to an external system.

Classification of prompts
Users' prompts are classified into a list of approved categories; the AI assistant must process or prohibit categories that it is not allowed to process. The classification is typically be carried out with a degree of certainty, which is why there is a risk that some prompts may still get through to the model. Classification is typically quick to perform and requires smaller and more specialised models than language models.

Security models
AI assistants should be considered as a system user in the organisation's system landscape and may therefore only access data and systems on behalf of authorised users. Access must be limited to relevant data. When AI assistants work in citizen-specific contexts, their func-tions and responsibilities must also be based on the access that the citizen has as a user.

A clear definition of tasks, setting of boundaries for areas of responsibility and access to available background information not only enhances precision but also increases confidence in the AI assistant.

3.6 Point 6 – Build structured quality assurance

Quality assurance is essential during both the development and implementation of secure and responsible AI assistants. Language models' responses are not deterministic. This means that their responses may vary, even if the same question is repeated, making the quality assurance process particularly challenging.

Effective quality assurance therefore requires a structured approach that includes a development phase, a pilot phase and continuous monitoring. In addition to traditional testing methods, such as functional tests and performance tests, methods specifically designed to address the AI assistants' patterns of behaviour are required so that they are not only technically well-functioning but also deliver value to the organisation in a responsible manner. As with other systems, standard tests such as functional tests or performance tests must still be maintained, especially in connection with the AI assistant's functional tools.

This point focuses on a number of methods specific to AI assistants that organisations should consider when performing quality assurance of AI assistants.

Business-driven design
An AI assistant should be designed in close collaboration between business experts and developers. Business experts establish criteria for how the AI assistant should respond in different scenarios. Developers translate the criteria into a technical design, which should include prompts, access to relevant knowledge sources and integrations with business systems. End users can be involved early in the process by testing different combinations of scenarios, prompts and background information on the model providers' websites.

If developers design the AI assistant alone, without obtaining knowledge from the business, the organisation risks that the AI assistant only function in a limited reality based on the developers' understanding. This can lead to the AI assistant being unable to address domain-specific challenges unknown to the developers and, in the worst case, the AI assistant making incorrect decisions or providing inadequate responses.

Red Teaming
An approach in which experts with technical and business insight into the AI assistant preventively identify the risks of harmful results when using the AI assistant. An example of this could be the AI assistant delivering promises that the organisation is unable to meet or sharing data that should not be shared. There are two approaches to red teaming:

- **Neutral:** The expert tests the functionality for which the AI assistant is designed based on expected use. Risks identified through the neutral approach are critical, as they may occur more fre-quently compared to risks identified through the hostile approach.

- **Hostile:** The expert actively tests the AI assistant to identify risks. Anything goes, and the expert must identify both probable and unlikely risks.

Pilot project
After an AI assistant has been developed, but before its broad deployment within the organisation occurs, a pilot is run with a group of business experts who test the AI assistant in their actual work. Experts provide regular feedback on the AI assistant to the developers so they can improve the quality of its results. Language models are non-deterministic and, for AI assistants accessed through chats in particular, there is a significant risk that the design will not anticipate all outcomes. Throughout the pilot project, a core group of superusers who can guide and support their colleagues in using the AI assistant is built up. They can help the organisation with the further deployment by supporting their colleagues. A pilot project can be run with external users such as customers, where an invited target group tests and provides feedback on the AI assistant. As the target group is typically not able to provide continuous active feedback, these projects will be more of a data collection process.

Control

To support ongoing quality management, the organisation should monitor the AI assistant's results while it is in production. Controls can be carried out using a variety of methods:

- **Manual control:** A developer or business expert manually reviews the AI assistant's responses to identify incorrect or ambiguous responses. Manual review can be used in every phase of an AI assistant's implementation, but should be used when rolling out the AI assistant in particular, in order to detect unforeseen risks. Manual control benefits from expert-driven insight and human intuition. However, the heavy workload and dependency on experts does limit the efficiency of manual control. Over time, manual control should be automated to facilitate quality assurance when modifying the AI assistant.

- **Rules-based control:** Control of the AI assistant's results based on definable rules. If the expected result includes a postal address, for instance, rules-based control can be used to ensure that the field is included and correct. A rules-based approach is transparent and allows the use of a percentage rate to measure quality. The controls require ongoing maintenance or additions as the functionality of the AI assistant expands and they cannot identify hidden or abstract uncertainties.

- **Model-based validation:** The AI assistant's results is validated before being presented to the users, typically using a language model or a different AI model. The performance of the AI assistant falls off because the user has to wait for the results to be validated. On the other hand, the risk of error, prompt manipulation or hallucination is reduced, and the model can identify unknown and abstract uncertainties. For certain critical business processes, precision is more important than latency or cost, which makes automatic validation a preferred solution. Model-based validation can also be used without user involvement as part of the testing, where a number of questions and responses are automatically processed by models to ensure the solution maintains quality before a new version is launched.

Structured documentation

It is critical that the AI assistant and quality assurance process are formalised and well documented. The following document types should be included in the documentation:

- **Functional design:** Describes the AI assistant's design-level responsibilities, role and purpose based on the AI assistant's use case.

- **User guide:** Describes how to use the AI assistant and its limitations.

- **Legal documentation:** Describes the data processing as well as the necessary risk assessments to be carried out based on the GDPR and AIA. See section 3.4 as regards the legal framework.

- **Technical design:** Describes the technical structure of the AI assistant, followed by prompts, internal functions and how it integrates with other systems.

- **Security documentation:** Describes security measures, including access control and procedures for handling security incidents. The documentation should also include an assessment of potential vulnerabilities and measures to mitigate these.

- **Test design:** Documents that the AI assistant acts appropriately in regard to requirements and design. Since language models are not deterministic, tests must focus on what the expected result is expected to contain.

- **Operational documentation:** Describes day-to-day operating procedures such as monitoring and routines for updating and maintaining the AI assistant.

The underlying language model used by the AI assistant must also be documented. The documentation should demonstrate how it is ensured that the model works with information made available in its prompt and how bias is addressed. If an existing model is used, its developers should have the necessary documentation.

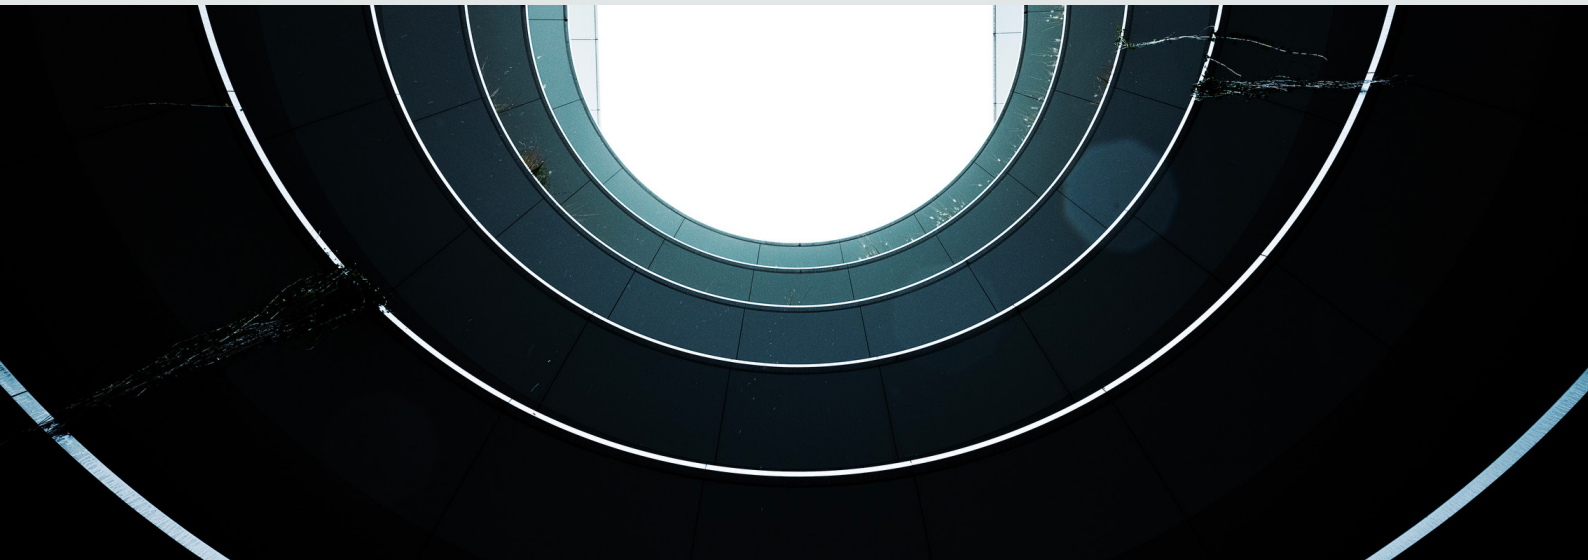

3.7 Point 7 – Measure and store relevant data on the use of the AI assistant

To ensure transparency and compliance with requirements, it is necessary to establish a detailed data audit and lineage. The goal is a setup that enables logging of all interactions and actions performed by AI assistants and provide full transparency in relation to data flow and the AI assistant's decision-making.

- **Data audit:** The process for logging and monitoring all relevant data activities. This makes it possible to document how data was used, who has had access and on which data the decisions are based.

- **Data lineage:** Mapping of the movement of data from origin to destination. Tracking provides insight into what data has been used and how it has been transformed and processed along the way. For example, when using RAG, it may be relevant to log what the RAG process introduced for a given prompt.

In the case of data auditing and logging of AI assistants' activities, it is important to distinguish between two types of logging:

- **General logging:** Used when the AI assistant does not have access to personal data. This type of logging focuses on anonymous, overarching interactions.

- **Context-specific logging:** Used when the AI assistant has access to personal data. Context-specific logging is more detailed and links activities to specific citizens.

In a general context, where the AI assistant does not have access to personal data, any citizens are processed anonymously. This typically includes logging of prompts, the AI assistant's responses and the context behind those responses, without linking the information to a specific citizen. On the other hand, if the AI assistant works with a citizen's personal data, context-specific logging is used. In this case, the logging must include information about the specific citizen whose information is processed.

This distinction between general and context-specific logging is not only important for traceability but also to demonstrate to supervisory authorities and other stakeholders that the AI assistant is operating in accordance with relevant legislation.

It is essential to have a centralised and secure solution in place on the platform to ensure the correct logging and handling of data. The solution should have strict access controls so that only authorised personnel can access and administer the log files. In addition, depending on what the AI assistant is used for, it may also be necessary to log data at a very detailed level. For many AI assistants, this means being able to produce reports that show how they are used right down to each individual conversation. This makes it possible to monitor and analyse how data moves and how decisions are made.

Storing and handling the large volumes of data requires planning. As the amount of data logged may be significant, there should be clear rules, based on the GDPR and AIA, on how long the data is stored. According to the AIA, data must be stored for a minimum of six months and thereafter for an appropriate period of time depending on need. It is also important to ensure that data is stored in a way that protects against unauthorised access while also allowing for its retrieval and analysis when necessary.

To ensure the right security level, dedicated resources must exist to address problems. This team must be able to respond quickly and effectively to any security breach or data issues that may occur. Incident management is essential for maintaining confidence in the AI assistants, both internally within the organisation and externally among customers and supervisory authorities.

To establish a detailed data audit setup, relevant information covering specific data points is recommended. These data points are presented in Table 3.1.

Data Point	Description	Justification
AI assistant	The specific AI assistant that is associated with the other data elements.	Traceability requires identification of which AI assistant has processed a prompt.
User	The users who have interacted with the AI assistant. Users include human users but also other systems or AI assistants.	To ensure traceability and adherence to data ownership and access control requirements.
Session	A parent data element that binds other elements together and records the start and end of the session.	Sessions enable a holistic understanding of interactions and incidents.
Prompt	The specific prompt sent to the AI assistant. Sensitive information is masked prior to logging.	Necessary to track, analyse and validate the AI assistant's behaviour and to ensure GDPR compliance.
Authorisations	The AI assistant's authorisations to complete relevant tasks.	To ensure that the AI assistant is restricted to performing actions that are relevant to it.
Data and source	The data that the AI assistant has used to generate its result.	Documentation of data sources is important for understanding and verifying the decision basis.
Result	The generated output produced by the AI assistant.	As output from language models is not deterministic, it is logged to enable subsequent analysis.
Classification	How the prompt is classified by the AI assistant and an explanation of the generated result. A prompt may undergo multiple classifications.	Classification provides insight into the decision process and ensures comprehensibility of the results, as well as enabling follow-up in the event of incorrect classification.
Language	The language in which the message is written.	The language affects the understanding and interpretation of prompts, which is relevant to the quality of the results.
Context	The context or prompt engineering that the AI assistant has been exposed to during production of the result.	Documentation of context is important for understanding and reproducing the AI assistant's behaviour.
Quality assurance	Procedures for quality assurance of data and results. If the AI assistant performs quality assurance itself, this is logged.	Quality assurance is key to ensuring the accuracy and reliability of the AI assistant's output.
Dates and times	Date and time of all other data elements.	The timing of each action is critical for proper audit and analysis.

Table 3.1: Data points that are generally appropriate to log for AI assistants.

3.8 Point 8 – Plan organisational implementation and training

Organisations that implement AI assistants must train and prepare their employees for the changes that AI assistants bring. According to the AIA, there is a duty to ensure that employees, among others, are informed that they are interacting with AI assistants and that they receive the necessary training to use the AI assistants, taking into account their technical competence. Failure to adequately prepare the organisation and employees may lead to resistance to change.

This section presents four focus areas for this organisational implementation (Figure 3.5). This section does not include an exhaustive list of organisational implementation activities because the necessary initiatives depend on different parameters, such as the extent to which the organisation is affected, the specific target groups and the organisation's readiness for change and AI maturity.

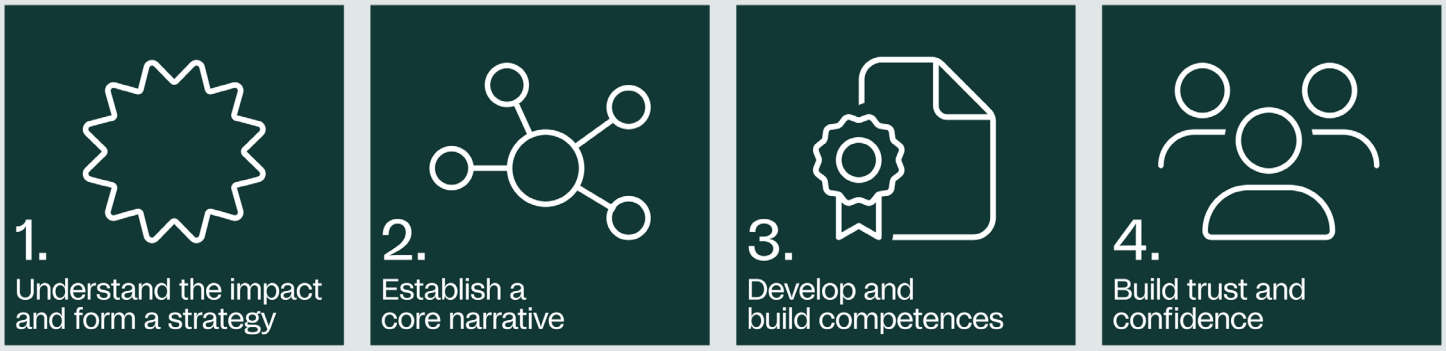

Figure 3.5: Organisational focus areas during implementation.

Focus area 1: Understand the impact and form a strategy
It is essential to understand how the organisation is affected by the introduction of an AI assistant. This can be achieved by examining the procedures and processes that is impacted and also by analysing which employee groups is affected. The implementation of the AI assistant affects not only the employees who uses it and the IT colleagues who maintains it but also the citizens who interacts with the AI assistant.

Through process, impact and target group analyses, the organisation can identify how to change the way it works and how processes and working procedures is affected. This helps to identify the need for new working methods, knowledge and skills. Once the extent of transformation is clear, the organisation must establish an implementation strategy and plan. This plan should ensure a smooth transition so the organisation is ready to use the AI assistant upon launch. It must take into consideration focus areas, such as education, training and communication, all of which are necessary to prepare the employees to use the AI assistant.

Focus area 2: Establish a core narrative
The next focus area is to establish a clear, structured and transparent narrative on the implementation of the AI assistant. A clear narrative helps communicate essential communication messages such as:

- **Why:** Create an overall understanding of why you want to implement the AI assistant and the value it will bring to your organisation. For example, this could be increasing the efficiency of working procedures and the elimination of manual work.
- **When and how:** Provide information on the process that the organisation must go through, such as when and how employees will receive training in using the AI assistant.
- **What does this mean for the employee:** Create an overview of what implementing the AI assistant means for the employee and what the employee can expect.

The messages can be based on the use case defined in section 3.1. The narrative is used to ensure that the right messages about the AI assistant get attention. This creates awareness, buy-in and support within the organisation. Failure to present the narrative correctly risks causing concern and confusion about the change, which may lead to resistance. For example, the employees who will use the AI assistant may be concerned that the AI assistant is going to replace them, thereby making employees resistant to the change. Instead, employees should view the AI assistant as a virtual colleague who performs a boring job, freeing up their time for other tasks.

Focus area 3: Develop and build competences
Training on how to use the AI assistant is key to successful implementation, and the establishment of training ensures that employees develop skills to use the AI assistant and understand their new working procedures. This applies to both the employees who will use the AI assistant as well as the IT colleagues who will maintain the AI assistant.

As part of this training, it is important that employees be introduced to how AI works. This creates insight into how AI assistants can add value to the organisation. It may be beneficial to use superusers at this point. Superusers are a group of specially selected employees who are trained to teach the rest of the organisation and be advocates for change. A network of superusers across the organisation contributes to a strong confidence in and sense of community around the AI assistant in the organisation and helps increase adoption of the assistant.

Practical training with specific exercises is recommended to build skills in using the AI assistant. This gives employees experience in using the AI assistant and knowledge of how it operates. However, the training method depends on the complexity of the AI assistant given that some tasks might require more interaction with the AI assistant than others. Employees should also have a general understanding of the information that the AI assistant draws on so that they are aware of the foundation of the responses it provides. If the AI assistant's results are used for further development, it may be relevant to provide training in data quality as well as on how new data affects the AI assistant's results. A better understanding of the AI assistant's underlying data contributes to accountability and transparency.

Focus area 4: Build trust and confidence
The last focus area is to build trust among the users of the AI assistant. These may be internal users, such as the organisation's employees, or external users, such as the organisation's customers.

During implementation, it is important for employees to feel comfortable about using both the AI assistant and the results that the AI assistant generates. Otherwise, employees will not be motivated to use the new tool and will instead continue to carry out their tasks as normal. For this reason, the core narrative and training need to focus on where data comes from and how it is processed. Likewise, employees must learn to quality assure the AI assistant's results. For example, employees must learn to decode when the AI assistant is hallucinating. This could include deploying control questions where the employee knows what the response should be. This helps employees gain a better understanding of how the AI assistant works so that they can feel more comfortable using the AI assistant.

To build customer trust, it may be beneficial to explain to customers that the assistance helps improve customer service by using help text where the AI assistant is used. The AI assistant should also provide personal service to customers, such as clarifying personal details to the extent that this is legal. This way, it is possible to increase customer confidence in using the AI assistant.

3.9 Point 9 – Establish follow-up and support structures

The digital transformation does not end at go-live. The organisation must continue to ensure that the new working procedures are followed and keep building confidence in the AI assistant. Structures must therefore be established that ensure effective maintenance and continued lawfulness. The goal of the follow-up and support structures is to facilitate change within the organisation, improve the AI assistant and minimise risk by identifying problems as soon as they occur. The following tools can be used:

- **Experts:** Business and technical experts who advise and answer users' questions. Business experts are used for peer-to-peer training, while technical experts guide the development of the AI assistant. The first group of experts can be created from a pilot project, as described in section 3.5.
- **Feedback:** Function where citizens can report faults. A good feedback mechanism helps identify specific areas in which the AI assistant is not responding properly or an area in which it should be able to perform but cannot. This makes it possible to adjust the model based on specific and precise examples. One simple feedback mechanism could be the option to mark whether the response is good or poor. However, what might appear to the user to be a poor response is not necessarily an incorrect response, which is why it may be appropriate to collect more information or suggestions for correct responses.

Monitoring: Monitoring logs and feedback so that reported challenges are handled before they affect citizens to a significant extent. Continuous monitoring of an AI assistant should be automated. Monitoring may be based on the data points introduced in section 3.2. There are two primary approaches to monitoring: Rules-based and intelligent.

- **Rules-based monitoring:** Rules are used to identify behaviour in the solution that does not occur as expected or in an optimal manner. For example, the quality of a response in RAG can be measured in terms of relevance and accuracy, provided that a prompt's distance is calculated, which may indicate that hallucination has occurred. Rules-based monitoring is only based on known problems.

- **Intelligent monitoring:** An AI assistant implemented to monitor the log and escalate if it finds anything inappropriate. This AI assistant can assess responses based on input and recognise attempts at security breaches, such as prompt injection attacks,[7] and upscale from there. If a cloud provider is used, it must be ensured that a data breach of the monitored AI assistant does not also result in a data breach of the monitoring AI assistant. The two systems should be separated so that a security breach in one system does not affect the other.

Automated escalation: Citizens seeking advice may feel that an AI assistant is not able to answer their questions. Where problems are dealt with internally through feedback and support, citizens can be placed in a frustrating situation. The AI assistant should therefore have built-in control mechanisms that can record if the same question is being asked repeatedly or it cannot find relevant information to assist the citizen in their process. This could be by providing a telephone number or guiding the citizen to a support chat with a human supporter. This ensures that citizens always receive the necessary help, even when the AI assistant cannot provide a satisfactory response.

Follow-up: Human follow-up of feedback and monitoring with the objective of making the AI assistant better ensure effective, continuous optimisation of the AI assistant's performance. Experts should regularly evaluate monitoring data and feedback reports, as well as implementing the necessary adjustments. Since follow-up can be time-consuming, the logs reviewed should be filtered in advance to reduce the workload. The combination of automatic and manual processes ensures that the AI assistant is constantly improved and remains reliable.

Maintenance of the knowledge base: To ensure that the AI assistant remains accurate and up to date, a process should be established for ongoing maintenance of this knowledge base. Although some updates can be carried out automatically, it is necessary to update its sources when new legislation is introduced, for instance. Experts who maintain knowledge sources should work together with both legal advisers and technical teams to ensure that sources are up to date and accurate.

These structures – the combination of experts, feedback mechanisms, automated monitoring, escalation and human follow-up – ensure that AI assistants are continually optimised and maintain the highest standards of accuracy, compliance and reliability in their interactions.

[7] A technique in which an AI assistant is manipulated by inserting unexpected text into its input in order to change the assistant's behaviour or cause it to generate unwanted or incorrect responses [10].

4 Conclusion

This white paper has presented an approach for how organisations can develop and implement AI assistants in the best, most innovative and responsible way at a time when Europe is facing the greatest challenges of recent years.

We have a unique opportunity to leverage technology to create more efficient, sustainable and robust societies, and although specific technologies and implementation methods are evolving over time, the fundamental considerations in the white paper will remain relevant as they are based on technology-agnostic good practice.

By applying this policy, organisations can develop a process for using AI assistants where legal requirements and risk are considered. This process not only supports implementation of current AI assistants but will also provide the basis for future AI projects.

A key advantage of following the white paper is that it supports increasing knowledge and skills within the organisation as regards to working with AI. This will make the organisation more experienced and better prepared to implement AI assistants within the organisation in the future.

Organisations are encouraged to consider an AI platform early in their AI journey, built on sound IT policies. By choosing a robust platform and implementing key monitoring and management components from the start, organisations can lay the foundation for more effective scaling and integration of future AI assistants.

While AI technology continues to evolve, organisations' commitment to responsible, ethical and effective implementation will remain unchanged. By building on the foundation set out in the white paper, organisations can confidently navigate the exciting landscape of AI assistants and exploit their full potential for the benefit of their operations and society as a whole.

The accompanying addenda present actual cases that provide concrete examples of how the principles contained in the white paper can be put into practice. The cases should be explored for further insight and inspiration.

5 References

[1] EU-U.S. Trade and Technology Council, EU-U.S. Terminology and Taxonomy for Artificial Intelligence – Second Edition. The Directorate-General for Communications Networks, Content and Technology, 2024. Accessed: 26 September 2024. [Online]. Available at: https://digital-strategy.ec.europa.eu/en/library/eu-us-terminology-and-taxonomy-artificial-intelligence-second-edition

[2] P. Béchard and O. M. Ayala, 'Reducing hallucination in structured outputs via Retrieval-Augmented Generation', April 2024, [Online]. Available at: http://arxiv.org/abs/2404.08189

[3] The Danish Data Protection Agency, 'Hvad er personoplysninger?' ['What is personal data?'] Accessed: 26 September 2024. [Online]. Available at: https://www.datatilsynet.dk/hvad-siger-reglerne/grundlaeggende-begreber/hvad-er-personoplysninger

[4] The Danish Data Protection Agency, the Ministry of Justice and the Agency for Digital Government, 'Behandlingssikkerhed – Databeskyttelse gennem design og standardindstillinger' ['Processing Security – Data Protection by Design and Default'], June 2018. [Online].
Available at: www.datatilsynet.dk.

[5] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance). EUR-Lex, 2016.

[6] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance). EUR-Lex, 2024. [Online]. Available at: http://data.europa.eu/eli/reg/2024/1689/oj

[7] C. and T. Directorate-General for Communications Networks, 'Retsakt om kunstig intelligens' ['Act on Artificial Intelligence']. Accessed: 26 September 2024. [Online].
Available at: https://digital-strategy.ec.europa.eu/da/policies/regulatory-framework-ai

[8] T. B. Brown et al., 'Language Models are Few-Shot Learners', 2020. [Online].
Available at: https://commoncrawl.org/the-data/

[9] European Union Network and Information Security Agency, 'Guidelines for SMEs on the security of personal data processing', p. 10, 2016, DOI: 10.2824/867415.

[10] CERT-EU Team, 'CERT-EU Security Guidance 23-002 Potential impact and risks of Generative AI in EUIBAs', May 2023. Accessed: 26 September 2024. [Online]. Available at: https://cert.europa.eu/publications/security-guidance/security-guidance-23-002---potential-impact-and-risks-of-generative-ai-in-euibas/

Appendix A

This appendix contains a legal checklist, the purpose of which is to further specify and operationalise point 4 of the white paper (section 3.4). The checklist thus addresses the legal environment and the questions that need to be investigated and resolved when working to clarify the legal framework for the AI assistant.

The checklist is for inspiration and guidance – and does not constitute an official list of answers. A specific legal assessment must be made in relation to the AI assistant in question, including compliance with other legal requirements that apply to your organisation. The checklist addresses the AIA and GDPR and is structured accordingly.

The checklist is essentially framed as questions, and the questions are based on underlying legal requirements set out in the AIA or GDPR. References are made to these underlying legal requirements where relevant.

Regardless of whether or not the AI assistant is covered by the Act, the checklist can be used to ensure that the AI assistant is developed/used in compliance with the principles of responsible AI, which are largely based on the same principles as the obligations under the AIA.

1 The AIA

1.1 Risk level and role
As stated in point 4 of the white paper, you must first clarify:
1. The assistant's risk level
2. Your role under the AIA

The checklist is structured in accordance with these clarifications, and the relevant questions and conditions are thus divided according to the risk level and role; see below. If the AI assistant is considered to be high risk, and you have the role of deployer, you must make yourself aware of the relevant point (specifically section 1.2.2). Please note, however, that you should familiarise yourself with the other requirements so that you are informed of these, including ensuring that any supplier of the AI assistant is compliant with the relevant requirements, e.g. as a provider.

1.2 Risk level: The AI assistant is assessed as "high risk"

1.2.1 Provider requirements – high risk
As a provider, you must meet a number of requirements if the AI assistant is considered to be a high-risk AI system; see AIA, Section 2, Requirements for high-risk AI systems. Thus, as a provider, you must assess and ensure compliance with these requirements when developing use of the AI assistant:

1.2.1.1 Overarching obligations (AIA, Article 16)
In general, you have the following obligations as the provider of an AI assistant that is considered a high-risk AI system:
a) ensure that the AI assistant complies with the requirements set out in Section 2 of the AIA; see further details under points 1.2.1.2–1.2.1.7 below;
b) indicate on the AI assistant or, where that is not possible, on its packaging or its accompanying documentation, as applicable, your name, registered trade name or registered trademark and the address at which you can be contacted;
c) have a quality management system in place which complies with Article 17 of the AIA;
d) keep the documentation referred to in Article 18 of the AIA;

e) when under your control, keep the logs automatically generated by the AI assistant as referred to in Article 19 of the AIA;
f) ensure that the AI assistant undergoes the relevant conformity assessment procedure as referred to in Article 43 of the AIA, prior to its being placed on the market or put into service;
g) draw up an EU declaration of conformity in accordance with Article 47 of the AIA;
h) affix the CE marking to the AI assistant in accordance with Article 48 of the AIA;
i) comply with the registration obligations referred to in Article 49(1) of the AIA – is the AI assistant registered in the EU database?
j) take the necessary corrective actions and provide information as required in Article 20 of the AIA;
k) upon a reasoned request of a national competent authority, demonstrate the conformity of the AI assistant with the requirements set out in Section 2 of the AIA;
l) ensure that the AI assistant complies with accessibility requirements in accordance with Directives (EU) 2016/2102 and (EU) 2019/882.

1.2.1.2 Risk management system (AIA, Article 9)

1. Has a risk management system been established and implemented that is documented and maintained? (Article 9, no 1)
2. Does it cover the lifecycle of the AI assistant?
3. Have you mapped and analysed known and foreseeable risks associated with the AI assistant, including:
    I. estimating and evaluating the known and foreseeable risks that may arise when using the AI assistant,
    II. evaluating other risks that may possibly arise based on the analysis of data gathered from the post-market monitoring AI assistant.
4. Have you introduced appropriate risk management measures?

1.2.1.3 Data and data governance (AIA, Article 10)

1. Have you established governance and management procedures for ensuring the data quality of training, validation and testing data sets. For following data protection by design and data protection by default, see [4].
2. Have you ensured that training, validation and testing data sets are subject to such processes, in particular:
a) relevant design choices;
b) data collection processes and the origin of data and, in the case of personal data, the original purpose of the data collection;
c) relevant data-preparation processing operations, such as annotation, labelling, cleaning, updating, enrichment and aggregation;
d) the formulation of assumptions, in particular with respect to the information that the data are supposed to measure and represent;
e) an assessment of the accessibility, quantity and suitability of the data sets that are needed; f) examination in view of possible biases that are likely to affect the health and safety of persons, have a negative impact on fundamental rights or lead to discrimination prohibited under Union law, especially where data outputs influence inputs for future operations;
g) appropriate measures to detect, prevent and mitigate possible biases identified according to point (f);
h) the identification of relevant data gaps or shortcomings that prevent compliance with this Regulation, and how those gaps and shortcomings can be addressed.

1.2.1.4 Technical documentation (AIA, Article 11)

1. Have you produced all relevant technical documentation? See the documentation requirements in Annex IV to the AIA.
2. Have you ensured processes for continuous updating of the technical documentation?
3. Have you ensured that the documentation can be stored in accordance with Article 18 of the AIA?

1.2.1.5 Transparency, logging and provision of information (AIA, Articles 12–13)

1. Have you introduced measures to ensure transparency? Have you developed instructions for use? Have you introduced mechanisms to inform employees/users of the reasons and criteria for the AI assistant's results? Is this communicated in a clear and understandable way to the target group? Have you established processes to take into consideration user feedback and used this to customise the AI assistant? Have you communicated potential or perceived risks, such as biases? Depending on use, have you also considered communication and transparency to other target groups, third parties or the public?
2. Have you assessed whether the AI assistant's decisions, and thereby results, can be understood?
3. Have you introduced mechanisms to make it easier to review the AI assistant's output for employees/users, e.g. by ensuring traceability and recording of the AI assistant's processes and results?
4. Have you ensured that events ('logs') are automatically recorded in the AI assistant during operation? Have you used recognised standards or common specifications for logging?
5. Have you ensured that it can be explained why the AI assistant made a particular choice that led to a certain result in a way that is understandable to all employees/users who would like an explanation?
6. Have you assessed to what extent the AI assistant's decision-making influences the organisation's decision-making processes?
7. Have you designed the AI assistant from the start so that it can be interpreted? Have you researched and attempted to use the simplest and most interpretation-friendly model for the relevant application? Have you assessed whether you can analyse your training and test data? Can you change and update these over time? Have you assessed whether you are able to investigate interpretation capability following model training and development, or whether you have to assess the internal workflow of the model?
8. Have you specified what the purpose of the AI assistant is, and who or what can benefit from the product/service? Have the use cases been specified and clearly communicated, including with consideration to alternative means of communication, to ensure that they are understandable and appropriate for the target group in question? Depending on use, have you taken into account human psychology and potential limitations, such as the risk of confusion, confirmation of bias or cognitive fatigue?

1.2.1.6 Human oversight (AIA, Article 14)

1. Does the AI assistant interact with employees' and users' decision-making (e.g. recommended actions or decisions to be made, indicating options)? In this case, is there any risk that the AI assistant will interfere with human autonomy by intervening in the user's/employee's decision-making process in an unintended manner? Have you considered whether the AI assistant should notify users/employees that a decision, content, consultancy or conclusion is the result of an algorithmic decision?
2. If the AI assistant is a 'chatbot' or conversation system, are human users made aware that they are interacting with a non-human agent?
3. If the AI assistant is implemented in a working process, have you considered the distribution of tasks between the AI assistant and employees to ensure meaningful interactions and appropriate human control? Have you ensured that human oversight is performed by sufficiently competent employees?
4. Does the AI assistant improve or enhance human skills? Have you taken security measures to prevent excessive confidence in or dependency on the AI assistant in working processes?
5. Have you considered what level of human control is appropriate for the specific AI assistant and the specific application? Can you describe the level of human control or involvement, if this is relevant? Who has human control, and when and with what tools can human intervention take place? Have you introduced mechanisms and measures to ensure such potential human control or to ensure that decisions are made with human responsibility?
6. Have you taken steps to enable audit and correction of issues related to management of AI autonomy?
7. In the case of a self-learning or autonomous AI assistant or use case, have you introduced more specific controls and supervisory mechanisms?
8. What type of detection and response mechanisms have you established to assess whether anything could go wrong?

9. Have you established a ‘stop button’ or procedure that can abort an operation in a secure manner, if this becomes necessary? Does this procedure interrupt the process in whole or in part, or does it leave the control to a human?

1.2.1.7 Accuracy, robustness and cybersecurity (AIA, Article 15)

1. Have you ensured an appropriate level of accuracy for the AI assistant and stated this and the relevant parameters in the instructions for use?
2. Have you ensured that the AI assistant is resilient regarding any errors, faults or inconsistencies that may occur in the system or environment in which the system operates, in particular due to the system’s interaction with natural persons or other systems.
3. Have you ensured the robustness of the AI assistant through technical redundancy solutions, which may include backup or fail-safe plans?
4. Have you ensured that the AI assistant is resilient against attempts by unauthorised third parties to alter its use or performance by exploiting the AI assistant’s vulnerabilities?
5. Have you introduced appropriate technical solutions to address AI-specific vulnerabilities, such as measures to prevent and control attacks trying to manipulate the training data set (data poisoning), inputs designed to cause the model to make a mistake (adversarial examples) or model flaws.

1.2.1.8 Quality management system (AIA, Article 17)

1. Have you established a quality management system to ensure compliance with the AIA? Is the system documented in a systematic and orderly manner in the form of written policies, procedures and instructions?
2. Have you ensured that the documentation includes at least the following aspects:
a) a strategy for regulatory compliance, including compliance with conformity assessment procedures and procedures for the management of modifications to the high-risk AI system;
b) techniques, procedures and systematic actions to be used for the design, design control and design verification of the high-risk AI system;
c) techniques, procedures and systematic actions to be used for the development, quality control and quality assurance of the high-risk AI system;
d) examination, test and validation procedures to be carried out before, during and after the development of the high-risk AI system, and the frequency with which they have to be carried out;
e) technical specifications, including standards, to be applied and, where the relevant harmonised standards are not applied in full or do not cover all of the relevant requirements set out in Section 2 of the AIA, the means to be used to ensure that the high-risk AI system complies with those requirements;
f) systems and procedures for data management, including data acquisition, data collection, data analysis, data labelling, data storage, data filtration, data mining, data aggregation, data retention and any other operation regarding the data that is performed before and for the purpose of the placing on the market or the putting into service of high-risk AI systems;
g) the risk management system referred to in Article 9 of the AIA;
h) the setting-up, implementation and maintenance of a post-market monitoring system, in accordance with Article 72 of the AIA;
i) procedures relating to the reporting of a serious incident in accordance with Article 73 of the AIA
j) the handling of communication with national competent authorities, other relevant authorities, including those providing or supporting the access to data, notified bodies, other operators, customers or other interested parties;
k) systems and procedures for record-keeping of all relevant documentation and all information;
l) resource management, including security-of-supply related measures;
m) an accountability framework setting out the responsibilities of the management and other staff.

1.2.2 Requirements for deployers – high risk (AIA, Articles 26–27)

As a deployer, you must meet a number of requirements if the AI assistant is deemed to be a high-risk AI system. As a deployer, you must therefore assess and ensure compliance with the requirements set out below when developing the use of the AI assistant. However, as noted from the outset, you should also familiarise yourself with the requirements that apply to providers.

1. Have you ensured that the high-risk AI system is used in accordance with the instructions for use?
2. Have you arranged human oversight of the AI assistant by someone with sufficient competence, training and authority, as well as the necessary support?
3. Have you ensured relevant and representative input data in view of the purpose of the AI assistant?
4. Have you ensured that the operation of the AI assistant is monitored on the basis of the instructions for use?
5. Have you ensured that logs are kept if they are under the deployer's control?
6. Have you registered the AI assistant in an EU database (Article 49)?
7. Have you prepared a data protection impact assessment (DPIA) based on the information contained in the instructions for use?
8. Have you prepared a fundamental rights impact assessment (FRIA) in accordance with Article 27 of the AIA? This includes assessing the wider social impact of the AI assistant's use beyond the individual user, such as potentially indirectly affected stakeholders. Please note that the fundamental rights impact assessment is only relevant to 'deployers that are bodies governed by public law, or are private entities providing public services, and deployers of high-risk AI systems referred to in points 5(b) and (c) of Annex III'; see Article 27. In other words,

– AI systems intended to be used to evaluate the creditworthiness of natural persons or to establish their credit score, with the exception of AI systems used for the purpose of detecting financial fraud;
– AI systems intended to be used for risk assessment and pricing in relation to natural persons in the case of life and health insurance.

2 GDPR compliance

You must assess compliance with the following general legal requirements of data protection for processing personal data when developing and using the AI assistant. Please note that the list of data protection rules and risks mentioned is not exhaustive and varies from project to project, but it does provide a solid foundation for compliance with data protection law.

2.1.1.1 Overview of the processing of personal data and your role

1. Do you have an overview of the data processing, including what personal data is processed about which individuals for what purposes throughout the lifecycle of the solution?
2. Have you clarified your data protection legislation role, including the data processing for which you are the data controller?

2.1.1.2 Managing data protection rules throughout the lifecycle of the solution

1. Have you taken measures – in a documentable manner – to incorporate data protection rules throughout the entire lifecycle of the AI assistant, i.e. right from the design phase to monitoring the solution in operation?
2. Have you addressed whether you comply with the fundamental principles of, among others, lawfulness, fairness and transparency, purpose limitation, data minimisation and data quality in Article 5 of the GDPR?
3. Have you identified a basis for processing (legal basis) for the processing of the personal data in the development/operation of the AI assistant?

2.1.1.3 Duty of disclosure and the rights of data subjects

1. Have you ensured that the data subjects are informed of the processing of personal data during both the development and operation of the AI assistant?
2. Will you inform the data subjects if profiling is used, including its consequences?
3. If automated individual decision-making is used, will the data subjects be informed of this and will they receive at a minimum meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subjects?
4. If automated individual decision-making is used, has a clear legal basis for been identified for this, and can you manage the rights of the data subjects, including the right to obtain human intervention, where appropriate?
5. Are there procedures for managing the other rights of data subjects, including right of access, rectification and objection as well as the right to erasure?

2.1.1.4 Security, impact assessment and effective management of risks to the data subjects

1. Based on a risk assessment of the security of the processing, have you identified risks relating to the data subjects' rights and freedoms, and implemented appropriate organisational and technical security measures throughout the lifecycle of the solution to address these risks?
2. Have you considered whether a data protection impact assessment (DPIA) will be conducted and, if so, have you conducted the assessment with the involvement of a data protection officer (DPO)?
3. Have you produced guidelines and given instructions to employees as how they should use the AI assistant correctly and securely to ensure that they are aware of the AI assistant's purpose and limitations, legal use and the risks involved?

2.1.1.5 Transfers to third countries and sharing of personal data with the supplier

1. Have you identified all possible transfers of personal data to third countries and assessed whether these transfers are lawful by preparing a transfer impact assessment (TIA) for this?
2. If you are using a Software as a Service solution, have you then reviewed the supplier's contract terms and data processing agreement to investigate whether, when using the solution, you disclose personal data to the solution supplier for the supplier's own purposes, e.g. training the solution, product development, marketing, etc. In this case, have you assessed whether this disclosure of personal data is lawful?


Netcompany
Microsoft
Topdanmark
JYSKE BANK
Poul Schmith
KAMMERADVOKATEN
Digitaliserings-
styrelsen
atp
ERHVERVSSTYRELSEN
KOMB:T
PFA
TRIFORK.
Dubex:
spar nord