Holistic AI White Paper

Requirements for High-Risk AI Systems

January 2023

Holistic AI

holisticai.com

Contents

Key takeaways

- The EU AI Act provides a set of requirements concerning different steps of the lifecycle of an AI system, starting from the data collection period to post-market monitoring for high-risk AI systems.

- Providers are the primary actors obligated to ensure compliance with these requirements, whereas some requirements are imposed on users as well.

- Requirements directly related to AI systems are provided under Chapter 2. However, there are additional requirements for providers under other provisions.

- Requirements are drafted broadly; although some provisions seemingly contain detailed instructions, the methodology of compliance may not be clear for all AI systems, given the flexible definition and wide range of AI systems covered by the Act.

- The content and scope of each requirement should be examined and determined after reviewing a given AI system and taking the "generally acknowledged state of the art" into account.

- Not all requirements are imposed on all high-risk AI systems. The determination as to which requirements apply to a given AI system should be made prior to proceeding with the compliance.

Introduction

The mass adoption of artificial intelligence (AI) is ubiquitous across sectors. While this can have many benefits, the use of these systems can pose novel risks if left unchecked. Precipitated by high profile cases of harm, such as the glitch in Knight Capital's trading algorithm where $440 million USD was lost in 30 minutes; the State of Michigan reached a $20 million USD settlement with Michigan residents wrongly accused of fraud by an automated system used by the state; and the Dutch Tax Authority scandal where 10,000s of lives were ruined after an algorithm was used to detect suspected benefits fraud, there has been increased industry, public and regulatory concern, with an impetus to manage the risk.

The European Union (EU) is leading the way with its proposed AI Act, which seeks to ensure that AI systems placed on the EU market are safe and do not pose a risk to the fundamental rights of citizens. Once adopted, we expect the Act to become a de facto global standard, similar to the GDPR in the data protection space. The EU AI Act (EU AIA) proposes a "risk-based approach" for regulating AI systems, where systems are classed as having (1) low or minimal risk, (2) limited risk, (3) high-risk, or (4) unacceptable risk. Systems classed as 'unacceptable risk' are banned and the risk management obligations vary depending on which of three other classifications apply.

High Risk Systems

Systems that can have a significant impact on the life chances of a user. There are 8 types of systems that fall into this category.

Limited or Minimal Risk Systems

Systems that:
- Interact with humans
- Detect humans or determine a person's categorization based on biometric data
- Produce manipulated content

Low Risk Systems

Includes spam filters or AI-enabled video games, and comprise the majority of the systems currntly being used on the market.

Systems classed as ‘high-risk’ must meet the requirements in Chapter 2 of the AI Act. These broad, stringent requirements related to the development, deployment, and usage of a high-risk AI system’s lifecycle. Important to note, Chapter 2 is not the only source of obligations relating to high-risk systems. Specifically, Chapter 2 sets out the “legal requirements for high-risk AI systems in relation to data and data governance, documentation and recording keeping, transparency and provision of information to users, human oversight, robustness, accuracy and security.”

Examples of key requirements elsewhere in the AI Act include:

- Chapter 3 obligations on other actors in the AI value chain (e.g., importers, distributors, authorised representatives, and so on).

- Chapters 4 and 5 detailing obligations for notified bodies and conformity assessments. These provisions broadly align with the EU’s product safety mechanisms.

This paper is not intended as an exhaustive guide to the Act. Instead, we focus primarily on the Chapter 2 requirements, aiming to explain them and to highlight an apparent gap between what is necessary for compliance and what will be considered sufficient to be not liable. We begin with a brief overview of Chapter 2 and other pertinent articles from across the Act referencing the General Approach version of the legislation, before outlining and analysing key provisions of Chapter 2’s articles.

Overview of chapter 2 of the EU AI Act & other pertinent provisions

Requirements for high-risk AI systems are provided under Chapter 2 between Articles 8-15. It must be emphasised that these are requirements ***directly*** related to the development, deployment, and usage of AI systems. It is on ***providers*** of high-risk AI systems to ensure that these systems are compliant with requirements.

Providers

‘provider’ means a natural or legal person, public authority, agency or other body that develops an AI system or that has an AI system developed and places that system on the market or puts it into service under its own name or trademark, whether for payment or free of charge.

Users

The notion of ‘user’ referred to in this Regulation should be interpreted as any natural or legal person, including a public authority, agency or other body, using an AI system under whose authority the system is used. Depending on the type of AI system, the use of the system may affect persons other than the user.

However, ensuring compliance the requirements set forth in these articles is not the sole obligation of providers; they must also ensure that high-risk AI systems undergo a conformity assessment to determine whether all of the requirements are met prior to placing them on the market or putting them into service, pursuant to Article (16)(e). This is mainly a self-assessment, except for the cases where an external assessment by a notified body is required. Consequently, Article 26(1)(a) imposes an obligation on importers of these systems to verify whether the relevant conformity assessment has been carried out by the provider. Additionally, Article 51 provides a registration requirement for some high-risk AI systems and their providers, and providers of high-risk AI systems are required to establish a post-market monitoring system under Article 61. In the Czech Presidency's final draft, Article 51 was updated, where the obligation for high-risk providers to register on the EU database for high-risk AI systems has been extended to public body users (such as public authorities or agencies) except for law enforcement. Failing to comply with these obligations could trigger hefty **penalties and administrative fines** according to Articles 71 and 72, respectively.

Additional requirements imposed on providers are not entirely isolated or independent from the requirements provided under Chapter 2 but are included throughout the text. For example, Article 43 details the process for conformity assessments to enable compliance with the requirements provided under Chapter 2. Similarly, details of the registration and post-market monitoring requirement outlined in Article 61 contain frequent references to the components of Chapter 2 requirements.

Article 8 stipulates that high-risk AI systems shall comply with these requirements, considering the generally acknowledged state of the art, which refers to best practices in EU law. Recital 49 of the EU AI Act also stipulates that *"high-risk AI systems should perform consistently throughout their lifecycle and meet an appropriate level of accuracy, robustness and cybersecurity in accordance with the generally acknowledged state of the art."*

Despite this broad wording, the reference to the generally acknowledged state of the art was provided solely with respect to risk management measures to be initially implemented under Article 9 in the Commission Proposal. This reference moved to Article 8(1) with the French Presidency's Compromise Text, which transformed the generally acknowledged state of the art as an evaluation metric for all requirements. This can be considered an improvement as AI systems are not standard, and not all systems may be equally suitable for the full implementation of all requirements. For example, the state of the art may vary depending on the type of system and the context.

Following Article 8, specific requirements for high-risk AI systems are provided under seven groups from Article 9 to Article 15:

1	**Risk management systems**	Article 9
2	**Data and data governance**	Article 10
3	**Technical documentation**	Article 11
4	**Record-keeping**	Article 12
5	**Transparency and provision of information to users**	Article 13
6	**Human oversight**	Article 14
7	**Accuracy, robustness and cybersecurity**	Article 15

Chapter two requirements for high-risk systems

● Risk management systems

The first requirement is the establishment, implementation, documentation, and maintenance of a risk management system (“**RMS**”) for high-risk AI systems.

Article 9(1) requires the RMS to be:

i. continuous,
ii. iterative,
iii. planned and run throughout the entire lifecycle of the system, and
iv. subject to regular, systematic updates.

The establishment of an RMS is also a part of the quality management system to be established by the providers under Article 17, and its detailed description is part of the technical documentation, which is to be drawn up as per Article 11.

RMSs shall be comprised of three main steps:

1. **Identification of known and foreseeable risks:** The focus on ‘known and foreseeable’ risks is narrower than the previous proposal to require a much broader risk assessment. Providers should take the system’s intended purpose into account when considering whether a risk is ‘known and foreseeable’. The Act defines ‘intended purpose’ as *“the use for which an AI system is intended by the provider, including the specific context and conditions of use, as specified in the information supplied by the provider in the instructions for use, promotional or sales materials and statements, as well as in the technical documentation.”* Indeed, to determine whether a given risk is as such, it should be first identified and analysed. In other words, the practical implication of this addition will not, in fact, be during the identification or analysis phase, but during the determination and subsequent adoption of suitable measures.

2. **Evaluation of possible arising risks based on post-market monitoring:** Article 61 of the EU AI Act requires providers to establish a post-market monitoring system according to a post-market monitoring plan. The post-market monitoring plan is a part of the technical documentation under Annex IV. The template and the details of the content of this plan shall be determined by the Commission as per Article 61(3). As a step of RMS, in addition to known and foreseeable risks, other potential risks, if any, should be identified and analysed based on the data and analysis of this post-market monitoring system. To this end, a notable addition made by the French Presidency's Compromise Text states that risks referred to under Article 9(1) are only those which may be reasonably mitigated or eliminated through the development or design of the high-risk AI system or the provision of adequate technical information. Complete elimination is not required; *the aim is to reduce the so-called high-risks to an acceptable residual level.*

3. **Adoption of suitable measures:** After identifying and analysing known, foreseeable, and, if applicable, post-market monitoring-related risks, suitable risk management measures shall be adopted. These measures aim to reduce the risk associated with high-risk AI systems as much as possible and to an acceptable level. Measures shall be determined, taking all other requirements into account. There needs to be a balance.

When is risk management "sufficient" for the purposes of the Act?

Currently, there are no specific criteria to ensure a risk management system has identified risks *sufficiently* enough to absolve liability; rather Article 9 broadly describes what an RMS should include. What is sufficient will likely be derived from the technical standards that the European Council has tasked the European standards organisations (CEN/CENELEC) with and out of further implementation clarities that will become apparent once the Act goes into force. Additionally, Article 9 stipulates that it is necessary that an RMS identifies and analyses the known and foreseeably most likely risks to occur in the context of health, safety and fundamental rights. However, using broad language such as "most likely" and "foreseeable" makes it hard to know to what extent the EU will be receptive to providers interpreting this broadly.

In addition to the 3 key steps, Article 9 also identifies the most appropriate risk management measures:

- *Elimination or reduction of risks identified and evaluated pursuant to paragraph 2 as far as possible through adequate design and development of high-risk AI system.*

- *Where appropriate, implementation of adequate mitigation and control measures in relation to risks that cannot be eliminated.*

- *Provision of adequate information pursuant to Article 13, in particular as regards the risks referred to in paragraph 2, point (b) of this Article, and, where appropriate, training to users.*

This section of Article 9 also outlines necessary measures for RMS; however, broad wording makes it difficult to assess where the line of liability lies. For example, the first point – "as far as possible" subjectively denotes the extent to which providers must adjust design and development to eliminate or reduce risks. Although slightly clearer than the first, the second point also creates some ambiguity since it is necessary to identify risks and eliminate them (at least those which can be mitigated) but for those that cannot be eliminated, they must be mitigated, *where appropriate*. The line of liability itself will therefore be an iterative process and liability cannot be forecasted entirely; instead, an analysis of potential implementation challenges would serve well. Further, what is sufficient is also dependent to an extent on whether the high-risk system is *likely to be accessed by or have impact on persons under the age of 18* (Article 9 – (8)). Nevertheless, it is important to consider that amendments to the AI Act in parliament could make it clearer where the line exists to be absolved from liability. However, if this is not the case providers should be prepared to follow the text as closely as possible until implementation standards are communicated (see Recitals 60 & 61).

In the interim as providers are preparing for the AI Act, providers should keep the following in mind:

- The establishment of a post-market monitoring system (referenced in Article 9, outlined in Article 61) – as a part of providers' RMS, data from the post-market monitoring system should be analysed to predict/ mitigate the evolution of other possible arising risks.

- An RMS that adopts from a framework or is developed as a framework which is iterative throughout the AI lifecycle – work by Jonas Schuett suggests taking note from already developed standards such as those proposed by the National Institute of Standards and Technology (NIST), or ISO/IEC DIS 23894. This is supported by legislation text which refers to "the generally acknowledged state of the art" in reference to RMS.

- Focus on the term iterative as was used above – RMSs should have regular systematic updating integrated.

- The standard that should be used when assessing whether a provider has complied with their risk management requirements is currently unclear. We could judge an RMS on whether it meets the Article 8 criteria, or on whether it protects individuals from harm in practice. The wording in Recital 42 suggests the latter. Recital 42 clarifies that the aim is to "mitigate the risks from high-risk AI systems placed or otherwise put into service on the Union market for **users and affected persons**, certain mandatory requirements should apply, taking into account the intended purpose of the use of the system and according to the risk management system to be established by the provider." Judging an RMS on whether it protects users in practice seems to be in line with the overall aims of the AI Act. Further, the overall original premise of the AI Act as a means of protecting individual rights would suggest that for an RMS to have sufficiently identified risks the risk management framework should consider both organisational and individual obligations.

● Data and data governance

The second requirement is related to the quality of the datasets used to train, validate, and test models for high-risk AI systems. Under Article 3 of the EU AI Act, training data is defined as *"data used for training an AI system through fitting its learnable parameters"; validation data is defined as "data used for providing an evaluation of the trained AI system and for tuning its non-learnable parameters and its learning process, among other things, in order to prevent overfitting; whereas the validation dataset can be a separate dataset or part of the training dataset, either as a fixed or variable split"; and testing data is defined as "data used for providing an independent evaluation of the trained and validated AI system in order to confirm the expected performance of that system before its placing on the market or putting into service."*

Recital 44 of the EU AI Act refers to high data quality as an essential element for the performance of systems. Accordingly, the Act itself provides certain requirements, which are, in principle, set forth for all data sets. However, for high-risk AI systems that do not involve model training, Article 10(6) foresees that these requirements shall only apply to testing data sets.

Firstly, Article 10 of the EU AI Act stipulates that these data sets shall be subject to data governance and management practices. The Act does not provide either a definition or a definitive list of these practices. Instead, it provides a non-exhaustive list of elements to consider:

- *The relevant design choices,*
- *Data collection processes,*

- *Relevant data preparation processes operations, such as annotation, labelling, cleaning, enrichment and aggregation,*
- *The formulation of relevant assumptions, notably with respect to the information that the data are supposed to measure and represent,*
- *A prior assessment of the availability, quantity and suitability of the data sets that are needed,*
- *Examination in view of possible biases that are likely to affect health and safety of natural persons or lead to discrimination prohibited by Union law,*
- *The identification of any possible data gaps or shortcomings, and how those gaps and shortcomings can be addressed.*

These are broadly worded practices, and each shall be evaluated according to the nature of the given high-risk system as well as the generally acknowledged state of the art, as previously explained.

Secondly, there are qualitative requirements for data sets, which must be *(i) relevant, (ii) representative, (iii) free of errors, and (iv) complete*. Noting that it is challenging to ensure that a given data set is complete and free from error, Article 10 recommends that these requirements should be fulfilled to the greatest extent possible and refers to the considerations of proportionality, technical feasibility, and availability of data under Recital 44. As such, the requirements shall be evaluated taking the intended purpose of the high-risk AI system in question into account and the groups of people on which it is intended to be used.

For AI systems that use personal data, there is likely to be an overlap between the AI Act and GDPR, which sets out broad obligations on the collection and use of personal data. Indeed, the EU AI Act has made it clear through its provisions that AI governance will not be achieved nor managed in a silo; providers must also refer to best practices of data governance and the GDPR to ensure compliance. As mentioned above, Recitals 44, 45 and 51 provide insight into potential measures that can be implemented to ensure data used in AI has the features cited above.

Examples of such measures include:

- Embracing practices of data sharing
- Bias monitoring, detection, and correction (this could potentially be in the form of audits or assurance although this is not mandated by the Act)
- Integrating cybersecurity best practices into data governance frameworks/measures
- Referring to cross-disciplinary teams or team members when using data; demanding that data sources if not developed in-house come with some sort of assurance

Additionally, the requirement of completeness cannot be understood as a ground to exclude the implementation of risk management measures and privacy-preserving techniques. For example, a privacy preserving measure designed to reduce the precision in data or to exclude some values from a dataset should not be ruled out on the basis that it renders the data incomplete.

In addition, Article 10(5) allows providers of high-risk AI systems to process special categories of personal data referred to in Article 9(1) of GDPR to the extent that it is strictly necessary for the purposes of bias monitoring, detection, and correction. This processing must be subject to "appropriate safeguards for the fundamental rights and freedoms of natural persons".

Thirdly, data sets need to take *"the characteristics or elements that are particular to the specific geographical, behavioural or functional setting within which the high-risk AI system is intended to be used"* into account to the extent required by the intended purpose of the AI system. When fulfilled, this requirement also triggers a presumption of conformity as per Article 42(1) of the EU AI Act.

● Technical documentation

The third requirement for high-risk systems is related to technical documentation, which should be prepared before the system is placed on the market or put into service and kept up-to-date. Placing on the market of an AI system is defined as "the first making available of an AI system on the Union market" whereas putting into service thereof is defined as *"the supply of an AI system for first use directly to the user or for own use in the Union for its intended purpose"* under Article 3 of the EU AI Act.

The technical documentation should demonstrate that the high-risk AI system complies with all the requirements laid out for high-risk systems and contain all the necessary information in a clear and comprehensive form to assess this compliance. Such documentation shall be prepared by the provider of the high-risk AI system and fulfil the minimum information required as outlined in **Annex IV** of the EU AI Act, which could be amended by the European Commission pursuant to Article 11(3). For SMEs and start-ups, any equivalent documentation meeting the same objectives should be provided.

Where the high-risk AI system in question is related to a product, additional legal acts listed the system must comply with are listed in Annex II. For these systems, the technical documentation shall contain additional information that is required by those legal acts.

Technical documentation is the key reference text for the conformity assessment, along with the quality management system to be established by the provider. It shall be kept at the disposal of national competent authorities as referred to under Article 3(44) of the act until after ten years from placing on the market or putting into service of the AI system.

Record-keeping

Article 12 of the EU AI Act requires AI systems to technically allow for logging, or automatic recording of events, during the lifecycle of the system. These systems should support logging in a manner that enables:

(i) identification of situations that may result in the AI system presenting a risk at a national level within the meaning of Article 65(1) of the EU AI Act or in a substantial modification,
(ii) facilitation of the post-market monitoring referred to under Article 61 of the EU AI Act,
(iii) monitoring of the operation of high-risk AI systems as referred to under Article 29(4) of the EU AI Act.

In addition to these general aims, for AI systems used for biometric identification, the Act provides the minimum content to be logged:

(i) recording of the period of each use of the system with the start date-time and end date-time of each use,
(ii) the reference database against which input data has been checked by the system,
(iii) the input data for which the search has led to a match,
(iv) the identification of natural persons involved in the verification of the results, as referred to under Article 14(5) of the EU AI Act.

Such logs should be kept for six months, unless provided otherwise in the national law or the Union law, to the extent that the AI system generates logs that are under the control of users.

Transparency and provision of information to users

The transparency requirements outlined in Article 13 are associated with maintaining compliance with other requirements as well as enabling users to understand and use the system appropriately.

To inform users, the act also requires high-risk AI systems to be accompanied by instructions, which should be concise, clear, accessible, and comprehensible to the users. In addition, Article 13(3) governs the content of these instructions, which should specify the following:

> *(i) the identity and the contact details of the provider and, where applicable, of its authorised representative,*
> *(ii) the characteristics, capabilities and limitations of performance of the high-risk AI system, including:*
>> *(a) Its intended purpose, inclusive of the specific geographical, behavioural or functional setting within which the high-risk AI system is intended to be used,*
>> *(b) the level of accuracy, including its metrics, robustness and cybersecurity referred to in Article 15 against which the high-risk AI system has been tested and validated and which can be expected, and any known and foreseeable circumstances that may have an impact on that expected level of accuracy, robustness and cybersecurity,*
>> *(c) any known or foreseeable circumstance related to the use of the high-risk AI system in accordance with its intended purpose, which may lead to risks to the health and safety or fundamental rights referred to in Article 9(2),*
>> *(d) when appropriate, its behaviour regarding specific persons or groups of persons on which the system is intended to be used,*
>> *(e) when appropriate, specifications for the input data, or any other relevant information in terms of the training, validation and testing data sets used, taking into account the intended purpose of the AI system,*
>> *(f) when appropriate, description of the expected output of the system,*
>
> *(iii) the changes to the high-risk AI system and its performance which have been predetermined by the provider at the moment of the initial conformity assessment, if any,*
> *(iv) the human oversight measures referred to in Article 14, including the technical measures put in place to facilitate the interpretation of the outputs of AI systems by the users,*
> *(v) the computational and hardware resources needed, the expected lifetime of the high-risk AI system and any necessary maintenance and care measures, including their frequency, to ensure the proper functioning of that AI system, including as regards software updates,*
> *(vi) a description of the mechanism included within the AI system that allows users to properly collect, store and interpret the logs, where relevant.*

In addition to informing users, this information will also be informing data protection impact assessments conducted by users that are required to conduct one under GDPR.

Human oversight

According to Article 14 of the EU AI Act, high-risk AI systems should be designed and developed so that they can be effectively overseen by natural persons while these systems are being used. It must be pointed out that human oversight is not the same thing as the provision of information or transparency. This is about involving humans in an active capacity.

Human oversight can be built into the AI system by the provider, and/or the provider may identify appropriate measures to be implemented by the user. In any case, human oversight measures should enable the person to whom the oversight is assigned to do the following:

> (i) fully understand the capacities and limitations of the high-risk AI system and be able to duly monitor its operation, so that signs of anomalies, dysfunctions and unexpected performance can be detected and addressed as soon as possible;
> (ii) remain aware of the possible tendency of automatically relying or over-relying on the output produced by a high-risk AI system ('automation bias'), in particular for high-risk AI systems used to provide information or recommendations for decisions to be taken by natural persons;
> (iii) be able to correctly interpret the high-risk AI system's output, taking into account, in particular, the characteristics of the system and the interpretation tools and methods available;
> (iv) be able to decide, in any particular situation, not to use the high-risk AI system or otherwise disregard, override or reverse the output of the high-risk AI system;
> (v) be able to intervene in the operation of the high-risk AI system or interrupt the system through a "stop" button or a similar procedure.

In addition to this general requirement, for AI-based biometric identification systems, no action or decision can be taken based on the identification provided by the system unless it has been verified and confirmed by at least two natural persons. However, national security, defence, and military purposes are excluded from the scope of the Act to ensure that, in extenuating circumstances, these agencies can use remote real-time biometric information.

These human oversight obligations do not only apply to providers but also users of high-risk AI systems. Article 29(1a) sets forth that users will assign human oversight to natural persons with the necessary competence, training, and authority. This obligation is closely associated with and based on the instructions to be provided by the provider of the AI system as per Article 29(4).

Accuracy, robustness and cybersecurity

As outlined in Article 15 of the high-risk AI systems should be consistently accurate, robust, and secure throughout the lifecycle of the system. As outlined above, information regarding the accuracy of the model should be provided within the instructions required under Article 13, with this information representing how well the system does what it is intended to.

Robustness is associated with the quality of being resilient against errors as well as inconsistencies. Such inconsistencies may occur during the operation of the AI system or interaction thereof with other systems or people. What is important is that there should be measures to minimise the effects of such events and keep the systems as robust as possible. Accordingly, high-risk AI systems that continue learning after being placed on the market or put into service should be designed with measures to mitigate potential biases.

Cybersecurity refers to the requirement that high-risk AI systems be resilient against unwanted third-party involvement, such as unauthorised uses or exploitations of vulnerabilities. The act refers to data poisoning, adversarial examples, or model flaws as examples of vulnerabilities against which measures need to be taken.

Framing this requirement around the system’s lifecycle implies a need for regular monitoring. For example, a system capable of learning and adjusting its behaviour based on inputs may experience model drift. In other words, the performance of a model can change over time.

Holistic AI can support you in your path towards compliance with the EU AI Act

As the explanations above demonstrate, whether a high-risk AI system is compliant with the requirements provided under the EU AI Act is an important question with many practical implications for everyone involved with the lifecycle of an AI system, primarily providers. Further, high-risk AI systems are associated with various, broadly drafted requirements, the content of which merits a separate examination. Expected to come into force within the next two years, the EU AI Act and developments thereof are already shaping the industry practice along with the formation of new rules as well as standards. Taking steps to manage the compliance risks of your high-risk AI systems is the best way to get ahead of this upcoming regulation. At **Holistic AI**, we have a diverse team of experts in computer science, algorithms, auditing, law, and public policy who combine their expertise to make AI more ethical, legal and safeguarded against potential harms.

To learn more about how we can empower your enterprise to adopt and scale AI with confidence, **get in touch with us today at we@holisticai.com**.

Authored by **Ashyana-Jasmine Kachra**, Public Policy Associate at Holistic AI, and **Osman Gazi Güçlütürk**, Head of the IT Law Department at Boğaziçi University and Visiting Fellow at Yale Law School

GOT QUESTIONS OR WANT TO SCHEDULE A CHAT? CONTACT US AT

🌐 holisticai.com ✉ we@holisticai.com

Holistic AI

DISCLAIMER: The information provided in this document does not, and is not intended to, constitute legal advice; instead, all information, content, and materials available on this site are for general informational purposes only. Information in this document may not constitute the most up-to-date legal or other information. This document contains links to third-party websites. Such links are only for the convenience of the reader, user or browser; Holistic AI does not recommend or endorse the contents of third-party sites.

Readers of this document should contact their attorney to obtain advice with respect to any particular legal matter. No reader of this document should act or refrain from acting on the basis of information in this document without first seeking legal advice from counsel in the relevant jurisdiction. Only your individual attorney can provide assurances that the information contained herein – and your interpretation of it – is applicable or appropriate to your particular situation. Use of, and access to, this document or any of the links or resources contained within the document do not create an attorney-client relationship between the reader, authors, contributors, contributing law firms, or committee members and their respective employers.

The views expressed in this document are those of the individual authors writing in their individual capacities only – not those of their respective employers, Holistic AI, or committee/task force as a whole. All liability with respect to actions taken or not taken based on the contents of this document are hereby expressly disclaimed. The content on this posting is provided “as is;” no representations are made that the content is error-free.