Extended Outline: Proposed Zero Draft for a Standard on Documentation of AI Datasets and AI Models

NOTE TO REVIEWERS

NIST invites input on any aspect of this document. Input on the following topics is especially welcome, in order of importance; reviewers with limited time may wish to focus on earlier items in the list:

- **The templates in clause 6**, in particular: the extent to which the templates strike an appropriate balance of detail and flexibility; what portions of the templates should be considered required, recommended, or merely informative suggestions; the feasibility of integrating dataset and model templates; and how to incorporate testing, evaluation, verification, and validation (TEVV) descriptors into them.
- **Any of the other topics highlighted by call-out boxes in the text**, which focus on scoping, desired qualities for AI dataset and model documentation, and trade-offs and guidance for navigating them.
- **The details of the expanded templates in the appendices**, particularly ways to make the template more specific or actionable.
- **The usefulness and sufficiency of the background and guidance in Clauses 4 and 5**, both on their own and as motivation for Clause 6 (see also note on scoping below).

Input can be shared by email to ai-standards@nist.gov. NIST welcomes input via marked-up documents, bulleted lists of comments and concerns, reaction letters, or any other form that stakeholders find most convenient. Submissions, including attachments and other supporting materials, will become part of the public record and subject to public disclosure. Organizations are also welcome to host listening sessions in which they gather stakeholders to share feedback verbally with the agency.

NIST will consider input received **by October 17, 2025** for the initial public draft of the text; input received later will be considered for incorporation into subsequent iterations.

The document includes two different kinds of text, which are formatted differently:

1. Content that is expected to be expanded into draft text in the zero draft is presented in Roman font.

2. *Content explaining NIST's thinking, plans, or decision-making for the benefit of reviewers of this outline is presented in italics, or in call-out boxes like this one.*

Background on the Zero Drafts Project

In March 2025, NIST announced its AI Standards Zero Drafts project. In this pilot project, NIST is collecting input on topics with a science-backed body of work and using it to develop "zero drafts"—preliminary, stakeholder-driven drafts of standards that are as thorough as possible. These drafts will then be submitted into the private sector-led standardization process as proposals for voluntary consensus standards. The project aims to broaden participation in and accelerate the creation of standards, helping standards meet the AI community's needs and unleash AI innovation.

Based on community input, NIST selected two topics for the Zero Drafts pilot, one of which is documentation of AI datasets and models[1] for transparency between AI stakeholders. This outline proposes a direction and structure for the forthcoming zero draft on AI dataset and AI model documentation. Based on this paper and input received in response, NIST will propose concrete text for a zero draft, which the community will also be invited to provide input on.

The resulting draft will be submitted to INCITS/AI, the private sector-led committee that represents the United States in ISO/IEC JTC 1/SC 42. SC 42 is the subcommittee that focuses on cross-sectoral AI standards development within ISO and IEC (two prominent international standards developing organizations that collaborate closely through joint committees). Assuming that INCITS/AI proposes the draft as a new project for SC 42 and that SC 42 takes up the proposal, the future of the document will be up to the usual consensus processes of standards development. NIST does not expect to maintain the document further, and will have no greater influence over the resulting standard than a typical contributor to an ISO/IEC national body. Given the document's intended trajectory, it aims to match ISO/IEC's language conventions (e.g., it uses "guidance" to refer to recommendations or "should" statements.)

Table of Contents

[1] *The topic originally selected included documentation of systems as a possible part of the scope. NIST is proposing to narrow that scope to datasets and models, as discussed in the scoping note below.*

- Clause 4: Potential Outcomes of Public-facing Documentation Activities and Artifacts
- Clause 5: Guidance on Public-Facing Documentation Artifacts and Processes
- Clause 6: AI Dataset and AI Model Documentation Templates
- Appendices

NOTE ON SCOPING

There are many high-level considerations for AI documentation, such as interested parties, potential benefits, and organizational process practices, that shape guidance for what to include in documentation and how to specify those documentation fields. These considerations are conceptually prior to detailed prescriptions about fields to document. This makes it difficult to formulate guidance on fields without discussing the higher-level considerations. Accordingly, several major clauses have been dedicated to motivation- and process-oriented material. To avoid excessively broad scope, the draft will become more of an overarching guide to documentation than a detailed examination of each potential element of a documentation template.

At the same time, it is important for a standard on documentation to specify what should be included in documentation artifacts, and stakeholders have indicated that such specificity is part of what they seek from this zero draft. Accordingly, templates do make up one of the main proposed clauses, and much of the higher-level content is included based on what is needed to motivate, interpret, and apply the templates.

The templates aim to provide sufficient guidance to create uniformity among documentation artifacts by specifying documentation fields whose inclusion has broad consensus, while deferring much of the detailed guidance on how to specify each element of AI dataset or AI model documentation to future work. In keeping with common ISO/IEC patterns (e.g., 27001 and 27002), a standard based on this draft could be the first in a two-part series, where the second part provides the more detailed guidance unpacking the templates.

Although NIST received substantial interest in system-level documentation, such documentation would involve many additional considerations such as architectural choices, interactions between models, and different system goals and use cases. Documentation of such considerations has been less thoroughly studied and implemented, making AI system documentation less amenable to standardization at this time. When the state of the art allows, system documentation would be a natural future extension.

NIST welcomes views on the resulting scope and structure.

Introduction

ISO/IEC documents optionally include an introduction commenting on "the technical content of the document and…the reasons prompting its preparation." The zero draft will include such a section to explain what the document intends to accomplish, why it is scoped the way it is, and how it is structured.

- Objectives of the document:
    - Provide well-motivated, practical templates organizations can use for publicly documenting AI datasets and models in a comparable, consistent way.
    - Clarify the intent and thinking behind those templates to guide interpretation and use.
    - Provide actionable guidance on conducting dataset and model documentation activities in ways that address common challenges and trade-offs.
- The need for a standard on documenting AI datasets and models:
    - Consistency between actors' practices on what and how to document would facilitate development, interpretation, and comparison of documentation artifacts.
        - There are currently many research analyses, guidelines, and templates regarding AI data and model documentation that are not clearly aligned.
    - A standard can also help organizations, especially those without the resources to develop internal processes and templates from scratch, to adopt strong documentation practices.
        - Such practices ultimately support building trustworthy AI datasets and models (see 4.2) while accounting for the practical challenges inherent in designing AI documentation processes and artifacts (see 5.3 and 5.4).
- Scoping:
    - Three types of AI-related objects lend themselves to documentation:
        - Dataset documentation: Provides model-agnostic information about a dataset's characteristics (e.g., intended use, design, evaluation, sources, and limitations).
        - Model documentation: Focuses on the characteristics of the model, including how a specific model interacts with and is influenced by its training and evaluation data.
        - System documentation: Describes an entire AI system, including models and all the software around them that orchestrates, constrains, and connects them.

- This document focuses on AI dataset and model documentation, which are typically intertwined.
    - While datasets can be documented on their own, documenting a model typically entails sharing some information about the data used to build the model.
    - Conversely, when documenting a dataset to shed light on a model, it can be helpful to modify or add to what one might say about the dataset alone.
    - Thus, dataset and model documentation artifacts often need to share some structural elements, and a dataset documentation artifact might even be nested inside a model documentation artifact.
    - AI system documentation is often similarly intertwined with model documentation, but it is out of scope, as it has been less thoroughly studied.
    - The draft will focus on public-facing documentation.
    - Much of the content in this document is general enough to apply to settings beyond public documentation, such as organization-internal or business-to-business documentation.
    - It is possible that an organization will find it necessary to create documentation beyond what is covered in the guidance and templates below (e.g., documentation that includes privileged, non-public, or simply more detailed information) for internal use or for controlled sharing with customers, contractors, business partners, or oversight bodies.
- The draft aims to be domain-agnostic and flexible.
    - Artifacts and processes will necessarily vary based on factors such as requirements, objectives, audience, data types, and originator.
        - E.g., for a single dataset, a dataset provider might separately publish an overview for general audiences, a targeted report for security and privacy professionals, and a thorough report for technical implementers.
        - Specialization of this document (e.g., “profiles”) may be helpful for some of these settings.
- Documentation practices rely on processes and considerations that span multiple organizational functions, not all of which are in scope.
    - E.g., broader governance and operations processes for managing AI systems, while relevant, are discussed only insofar as they focus on documentation specifically.

            - Guidance on such processes can be found in ISO/IEC 42001 and other relevant standards.
        - Such processes can also support non-public documentation efforts, e.g. for internal needs, audits, or customer documentation, which are connected but out of scope.
- Structure:
    - The outline builds up to recommended templates for use in AI dataset and model documentation. To motivate these templates and clarify how to interpret and use them, the document first provides informative content and recommendations that inform the design of documentation artifacts and processes.
    - Specifically:
        - Clause 4 describes noteworthy aspects of downstream outcomes that can be affected by documentation choices. These can be thought of as positive or negative impacts, although the framing deliberately avoids this characterization, as discussed there.
        - Clause 5 offers recommendations on desirable qualities for documentation artifacts and processes, as well as trade-offs between them that organizations creating documentation will need to navigate. The recommendations are geared toward achieving better outcomes on the dimensions discussed in Clause 4.
        - Clause 6 provides the templates for documenting AI datasets and models. The templates aim to instantiate Clause 5's recommendations on artifacts and integrate well with Clause 5's recommendations on process.
        - The appendices provide further suggestions and examples.

Clause 1: Scope

Per ISO/IEC convention, this clause will consist of one or two short paragraphs specifying what the rest of the document will cover. It will clarify that the standard will provide sector-independent guidance on purposes of and approaches to public documentation on AI datasets and models, as well as shared templates for documentation artifacts to follow.

Clause 2: Normative References

Normative references will be added as appropriate to incorporate external definitions or requirements. References in Clause 2 are limited to existing ISO documents.

Clause 3: Terms and Definitions

Per the usual ISO/IEC document structure, Clause 3 will contain definitions of key terms that underlie the discussion of AI dataset and model documentation.

Definitions will be crafted to reflect the conceptual discussion in Clauses 4 and 5, and will be aligned as closely as possible with existing ISO/IEC definitions (when these cannot simply be imported via normative references).

Terms defined will include at least the following, listed with provisional definitions to be refined later in the process:

- **Field:** entry in a structured documentation template describing a particular characteristic, or set of characteristics, of the object of documentation

Clause 4: Potential Outcomes of Public-Facing Documentation Artifacts and Processes

This clause will describe ways in which different AI documentation practices and artifacts can lead to notable variations in downstream outcomes.

The clause will focus on dimensions of outcomes (e.g., "internal decision-making" and "trust") on which interested parties (i.e., stakeholders) would likely evaluate how beneficial or harmful an outcome is. These outcome dimensions thus motivate and guide the creation of public-facing documentation. Nonetheless, the outline deliberately describes each dimension in terms of a range of possible outcomes along that dimension (where some outcomes may be normatively better than others), rather than evaluating specific outcomes as positive or negative.[2] This framing provides several advantages:

1. *It simplifies the treatment of the fact that some outcomes' beneficiality or harmfulness can differ, or be perceived differently, depending on context, use case, and the interested party making the judgment. (For example, whether it is beneficial for a model to be easy for members of the public to build on can depend on the model's positioning in the marketplace.)*

---

[2] Another way to conceptualize the distinction between dimensions and outcomes is in terms of random variables in statistics and probability. Each "dimension"—trust, accountability, resources consumed, etc.—can be thought of as a random variable representing a particular aspect of the state of the world that results from the documentation decisions (and other factors). A specific outcome with respect to a dimension—e.g., high trust in a company or product—is analogous to the random variable for trust taking on a particular value. Clause 4 focuses on characterizing the variables themselves, including the potentially large space of possible values each can take on, rather than identifying specific adequate or inadequate values for each.

2. *It helps consolidate normative statements and judgments, which typically take the form of "should" statements, in later clauses.*
3. *It accounts for the fact that even for a particular context, use case, and interested party, outcomes often are not simply beneficial or harmful in a binary fashion; rather, they fall on a sliding scale of benefit or cost along each dimension. For example, reliable documentation might provide maximal benefits with respect to trust, but superficial documentation might still provide greater benefits than no documentation. Whether one considers the trust outcome from superficial documentation beneficial, adequate, or disappointing thus depends on the point of comparison. By focusing on identifying the key dimensions of the outcomes, the clause avoids having to evaluate aspects of an outcome as being positive or negative in an absolute sense, allowing for a wide range of better or worse ways outcomes could materialize along a given dimension.*

For ease of exposition, the dimensions are treated independently here, but the dimensions do interact with each other. Trade-offs between the documentation qualities that affect each dimension are discussed in Clause 5.4.

4.1: General

- The clause identifies outcome dimensions of potential interest to an AI dataset or model creator that can be heavily affected by decisions about what to document and how.
    - Outcomes can present positive or negative impacts for various interested parties. Along each dimension, there are often a range of possible outcomes with varying degrees of benefit or harm to a given interested party, with documentation activities playing a significant role in which of those outcomes materialize.
    - The effects of documentation on these outcomes are often modulated by the degree of transparency that the documentation gives stakeholders into the characteristics and history of the dataset or model and its creators.
        - Transparency functions as an intermediate outcome that leads in turn to the outcomes of interest.
        - Consideration of these outcomes and the trade-offs between them can thus drive decision-making about how to properly calibrate transparency and how to design the documentation practices that yield that transparency.
- When considering documentation outcomes and objectives, it is helpful to distinguish between outcomes of documentation artifacts (via the information in the artifacts being publicly available) and outcomes of the process of creating documentation (via organizational impacts, independent of the impacts of the resulting artifacts). This document considers both.

- For clarity of exposition, potential outcomes are organized into those affecting only interested parties internal to the organization and those where the interested parties could be external, though there are many other ways outcomes and interested parties could be categorized.
    - Relevant interested parties include any group, organization, or individual who engages in or is impacted by the design, development, management, and use of AI datasets and models, whether or not these parties produce AI documentation.
        - Examples: dataset/model developers, dataset/model users, organizations that host or aggregate datasets/models (e.g., on a website), other integrators, public sector entities, impacted parties, regulators, organizational users, and end users.
    - The categorization is based on which interested parties experience the most immediate impacts, though other interested parties may be impacted as well as a downstream result.
        - E.g., the model creator may be affected by a malicious user's activity.
- While these outcomes can all result from public documentation, they are not necessarily unique to public documentation as opposed to other forms of documentation.
- Dimensions can interact with each other.
    - E.g., information about evaluation results that is useful for suitability assessment could also be helpful to malicious actors.
    - Trade-offs between the qualities of documentation artifacts and processes that typically lead to more desirable outcomes are discussed in Clause 5.4.
- Recommendations for how to achieve desired outcomes on these dimensions are addressed in Clause 5.

4.2: Potential impacts on interested parties within the dataset or model provider organization

- Decisions about documentation, including about what to document or not document and how, can have many kinds of impacts on parties within the provider organization, including:
    - **Resources:** The amount of an organization's time, energy, finances, and other resources consumed by documentation practices is usually proportional to how extensive those practices are. The extensiveness of documentation can also affect the resources required for later documentation efforts or other decision-making.
    - **Internal collaboration:** Decisions about documentation processes can facilitate collaboration within and between technical teams and governance teams to a greater or lesser extent.

    - **Internal decision-making:** Documentation practices can facilitate deliberative, context-sensitive decision-making with reliable information to a greater or lesser extent.
    - **Release speed:** Different documentation process designs can lead to slower or faster dataset or model release cycles.
        - This is particularly significant for more agile release cycles that involve publishing dataset or model updates at close intervals.
    - **Lifecycle management:** Documentation practices can make it easier or harder for technical teams to develop, maintain, and update datasets and models.
    - **Accountability:** By releasing or withholding information about decisions and decision-making, documentation can impact the extent to which the dataset or model creator's decisions can be examined and evaluated by others, particularly by interested parties outside the organization.
        - More thorough documentation can make development and decision-making processes more traceable.
    - **Information confidentiality:** Documentation practices can reveal or protect proprietary information.
        - E.g., releasing information about data may allow competitors or attackers to draw conclusions about products or services in development.
    - **Legal obligations:** Documentation practices can contribute, to varying degrees, to the fulfillment of or failure to fulfill obligations under relevant laws, regulations, or contracts (e.g., transparency or reporting requirements or obligations regarding data privacy or security).

4.3: Impacts on interested parties outside of the dataset or model provider organization

- The outcomes below most commonly apply to parties external to the dataset or model provider organization, but they can also apply to teams within the organization, particularly those that were not involved in developing the dataset or model.
- Documentation artifacts and practices can affect a number of outcomes and activities for external parties, including:
    - **Suitability assessments:** The existence of documentation artifacts and their contents affect recipients' ability to assess quality and suitability for their purposes, including by including or omitting testing, evaluation, verification, and validation (TEVV) results. For example, documentation artifacts can, to a greater or lesser extent:
        - Communicate characteristics that may be important for suitability (e.g., dataset sourcing or model performance).

        - Explain technical requirements for use and implementation considerations.
        - Communicate the objectives, findings, and limitations of TEVV activities and metrics to enable assessment.
        - Facilitate reproducibility of experiments or TEVV activities.
        - Improve evaluation practices themselves (e.g., the process of documenting how evaluation activities cover the intended use cases can reveal additional scenarios or test cases to assess).
    - **Reuse and integration:** Decisions about documentation can impact the ability of external parties to efficiently locate models or datasets and build on them.
        - E.g., documentation artifacts with conventionalized content and structure can help external parties search repositories for datasets or models more easily than artifacts with less standardized content would.
            - While implementers will often need detailed information, the presence of easily searchable fields can enable them to identify candidates for further investigation.
    - **Trust:** External parties' willingness to use or build on datasets and models depends partly on the extent to which they feel equipped to understand, in a manner that suits their needs, how well-considered the processes and decisions behind the product are and how the product could affect various interested parties.
        - Depending on the visibility a documentation artifact offers, it can leave external parties with varying degrees of understanding (or assurance that the desired understanding is attainable). External parties can also interpret the degree of visibility as a signal regarding the provider's openness, and thus potentially as an indicator of trustworthiness.
        - Building trust can require documenting many process- or decision-oriented aspects beyond the datasets or models' technical characteristics.
            - E.g., decision-making procedures, data collection and recruitment procedures, post-deployment experimentation procedures, etc.
    - **Privacy preservation:** Documentation artifacts can include sensitive information about individuals whose information is present in a dataset. The extent to which artifacts include or exclude such information determines how well they protect these individuals' privacy.
    - **Malicious use or attacks:** Documentation can provide malicious users or attackers with information that helps them misuse or compromise the dataset or

model. The thoroughness and detail included in a documentation artifact can affect how likely this outcome is to occur.

Clause 5: Guidance on Public-Facing Documentation Artifacts and Processes

This clause will contain most of the draft's "should" statements, referred to as "guidance" in ISO/IEC parlance, about both documentation artifacts and documentation processes.

The structure is as follows:

- *5.1 gives an overview of the clause and covers general recommendations on how to apply the recommendations in the subclauses below.*
- *Desired qualities of documentation artifacts and processes are discussed first in isolation in 5.2 and 5.3, respectively.*
    - *"Artifacts" is construed slightly more broadly than just the file(s) containing the documentation: 5.2 also covers recommendations for closely related outputs of documentation processes, such as websites, request forms, and other methods of accessing documentation.*
- *Documentation work will also need to account for the abundant trade-offs between these desired properties; such trade-offs are discussed in 5.4.*

5.1: General

- The recommendations below are intended to drive the outcomes discussed in Clause 4 toward those that organizations and other interested parties typically consider preferable.
    - Explicit connections are made to these outcomes where they are not necessarily obvious.
- All of the recommendations in this clause should be calibrated based on the contextual objectives of the organization, including the audiences the documentation artifacts are meant to serve.
    - In particular, not all documentation processes are geared toward public documentation.
        - Organizations should expect internal documentation processes to be more thorough than what is ultimately released as public-facing documentation.

        - These more thorough internal processes will often be the same ones that support internal, business-to-business, business-to-regulator, or other documentation needs.
    - The artifacts and processes should be geared toward whatever kinds of documentation are needed.
        - The recommendations below concern public-facing documentation, but are not necessarily applicable to other kinds of documentation.

NIST particularly welcomes input on whether the terminology used to describe the qualities listed below is suitable for public-facing documentation artifacts and whether there are additional key qualities or elaboration on qualities that should be included in the zero draft text.

5.2: Guidance on Documentation Artifacts

- Documentation artifacts should be tailored to the needs, uses, domains, interested parties, and technical outcomes that organizational objectives call for the artifacts to serve.
    - Public-facing documentation will typically be less comprehensive than documentation for limited external audiences (e.g., customers, partners, and regulators) or for internal audiences (e.g., developers, implementers, compliance personnel), particularly with respect to sensitive information, such as trade secrets or personal information.
- Specific qualities that should typically be sought in public-facing documentation artifacts include:
    - **Comprehensibility:** The presentation of information should be clear and accessible for relevant audiences.
        - Particularly for public audiences, this implies making the content as plain and concise as feasible.
        - Less comprehensible documentation artifacts will tend to produce less desirable outcomes regarding accountability, trust, suitability assessment, and reuse and integration.
    - **Informativeness:** Artifacts should be as extensive and detailed as needed for the purpose, describing as much as the audience needs to understand about the AI dataset or model and the development process and decision-making behind it.
        - This includes reflecting information from across the lifecycle, as needed.

        - "Needs to understand" should be considered in terms of what information interested parties will likely want to act on or use (e.g., for assessing suitability).
        - Contact points or other redress mechanisms can be important information to provide to audiences.
        - Omitting information essential to the audience can compromise trust, reduce accountability, inhibit assessing suitability, and prevent the documentation from assisting in reuse.
    - **Correctness:** Artifacts should contain factual, technically precise, evidence-based information describing the datasets or models.
        - The rigor and precision of the information should be tailored to the needs and constraints of the intended documentation audience, which may imply variation across different documentation artifacts.
        - Inaccurate artifacts will fail to achieve most positive outcomes while wasting resources, and can even expose organizations to legal, reputational, or other risks.
    - **Judiciousness:** Artifacts should include only information that is relevant and useful for public audiences and appropriate to share publicly.
        - Artifacts should omit private, proprietary, or legally compromising information.
        - Including inappropriate information can compromise confidentiality, violate legal obligations, reduce trust, facilitate malicious use or attacks, or compromise privacy.
    - **Freshness:** For documentation artifacts presented as canonical or current, the information conveyed by documentation artifacts should be up to date.
        - Older versions of documentation might still be useful, and organizations may want to consider whether and how to store, make available, mark, and otherwise manage these documents.
        - Organizations should ensure that documentation is available for the most recent version of a dataset or model.
    - **Interoperability:** Documentation artifacts should follow established conventions for content, structure, and format, matching artifacts for other AI datasets and models.
        - Organizations should provide both machine-readable and human-readable artifacts, which may or may not be the same artifact.
            - Human-readable artifacts support interested parties' decision-making, particularly for parties without technical AI

backgrounds; machine-readable artifacts support technical accessibility, reusability, and analysis of the products.
        - Each type of artifact should be tailored to the functions it supports.
    - Where the information does not lend itself to consistent and well-specified formal structures, artifacts can still include machine-readable indications of which free text or other unstructured content is part of what broad category of information (intended use, evaluation results, etc.).
- Interoperable artifacts enhance the findability of dataset and model documentation (e.g., via indexed metadata) and make the products easier to compare, ultimately enhancing trust, and reuse and integration, suitability assessments, and accountability.
- Interoperable artifacts also support comprehensibility: those immersed in the field will find it easier and faster to make sense of an artifact in a familiar format.
- Interoperability can support continuous documentation processes (see 5.3) by allowing portions of the documentation process to be automated.

- **Findability and availability:** Artifacts and the infrastructure surrounding them (e.g., access portals) should make it easy to find and access the artifacts.
    - Artifacts should be published with pointers (e.g., hyperlinks in intuitive locations), storage locations, descriptions or metadata, and possibly accompanying materials (e.g., blog posts) that make them easy to locate.
    - Artifacts can be made available via a number of mechanisms that make access easier or more difficult (e.g., open download, registration required, or tiered access). Mechanisms that make access easier should be preferred, to the extent consistent with organizational objectives.
- **Maintainability:** Artifacts and their components should be easy to create, extend, update, and maintain across the AI lifecycle.
    - Maintainability includes modularity.
        - Information should be separated into distinct components, each of which could easily be replaced with an expanded or condensed version if desired.
        - In some artifacts, it can be desirable to omit certain components.
    - Maintainability facilitates other important artifact and process qualities, including correctness, freshness, and process manageability.

5.3: Guidance on Documentation Processes

- The process recommendations below are intended to create reliable documentation practices that yield artifacts meeting the objectives in 5.2 and that otherwise drive preferred outcomes from Clause 4. The recommendations aim to account for the many kinds and sources of input that documentation can draw on.
    - Effective organizational governance processes rely on robust documentation processes and vice versa.
    - Organizations should build processes to define their requirements and objectives vis-a-vis the interests and needs of different interested parties, decide on trade-offs, and maintain and continually improve their practices.
        - For example, depending on use case, relevant regulatory requirements, industry vertical, and third-party interests, organizations may have to weigh the interests of different parties in areas such as transparency, completeness, and timeliness.
- **Provide organizational support:** To the extent called for by organizational objectives, organizational leadership should support documentation by allocating necessary resources (e.g., time, budget, headcount) and setting up institutional incentives for prioritizing documentation (e.g., management recognition, compensation via KPIs in performance evaluation, guidelines or tools).
    - Resources and incentives should be tailored to the objectives, including desired outcomes from Clause 4.
- **Define objectives:** Organizations should define their objectives for documentation, and design the remainder of their documentation processes and the resulting artifacts accordingly.
    - They should also assess, continuously or periodically, the extent to which their documentation objectives are being met.
    - Objectives will likely involve many outcomes from Clause 4.
- **Keep processes manageable:** Documentation processes should be manageable to execute given available resources. Otherwise, the organization risks exhausting its resources and potentially compromising the documentation altogether or even the model or dataset.
    - To mitigate complexity of collaborating on documentation and integrating it with other parts of the dataset or model lifecycle:
        - **Consider incremental artifacts:** Consider periodically creating internal documentation artifacts, perhaps partial or unfinished ones, with records from specific milestones. This approach helps documenters focus on recording relevant information at a given stage of development of an AI

dataset/model, thus providing a trail of artifacts that can be further consolidated into a final public documentation artifact by others.
        - **Support documentarians:** Designate documentarians and encourage them to play an active role. They should state concerns, raise questions, seek clarifications, and facilitate an environment of deliberation about the dataset or model development process as a whole.
- **Provide guidelines:** Organizations should create and maintain internal guidelines for producing documentation artifacts that satisfy the organization's and interested parties' needs.
    - Guidelines might need to reflect the tiering of documentation: organizations might need layers of guidelines, only some of which apply to producing public-facing documentation.
- **Document continuously:** Documentation activities should be undertaken throughout the lifecycle of the dataset or model.
    - Organizations should typically avoid leaving documentation until the dataset or model is largely complete.
    - Rather, documentation should be created in earlier stages of the lifecycle (e.g., in the design and development stage, planned data sources and collection processes could be documented). Documentation should then be revised and fleshed out in later stages based on more detailed or up-to-date information (e.g., implementation details), findings/results to date, and transformations to the dataset or model (e.g., model retraining, data pipeline changes).
        - It is often advisable for SMEs to dynamically create and adapt summaries of their technical work as they are doing it, in coordination with documentarians.
        - Continuous documentation could include documenting post-deployment impacts or experimental results.
    - This is necessary to keep public artifacts up-to-date, and helps to create traceability of the decision-making and design choices that facilitates adequate informativeness.
- **Distribute documentation work:** Documentation work within an organization should be distributed and collaborative.
    - Organizations should typically avoid assigning documentation responsibilities to a single person. Rather, organizations should provide for multiple functions that will typically be served by multiple people:
        - SMEs should provide asynchronous technical accounts of their own work developing specific dataset and model subcomponents (e.g. data collection, annotation, versioning, metrics, etc.).

        - Documentarians should oversee and keep track of documentation processes, coordinating between other documenters.
        - Organizations should establish guidelines and processes to support SMEs to document relevant information that support the documentation artifacts while minimizing costs.
    - Artifacts produced by distributed, collaborative processes will tend to be more informative and modular than those created by a single person, and the process will likely be more manageable and continuous. A distributed documentation process is also more likely to produce artifact-independent benefits such as facilitating internal collaboration and deliberative decision-making.
- **Incorporate audience input:** To the extent needed to achieve organizational objectives, documentation processes should allow for and incorporate feedback and other input from the intended audiences, potentially up to and including co-design of the documentation artifacts.
    - While organizations may not be able to directly survey these audiences, they can rely on other sources of information, such as surveys, relevant standards, industry publications, and tabletop exercises in which the organization considers what the identified interested parties require.
    - Factors that can shape the involvement called for by organizational objectives include mission, sensitivity of information, and redress or recourse needs.
    - Documentation processes are meant mainly to produce artifacts that help the intended audiences. Incorporating audience input can help to understand audience needs.
        - Audience input can help the provider to more quickly identify if existing artifacts are failing to exhibit desired qualities and/or to provide the intended benefits.
            - E.g., audience feedback can notify an organization if artifacts are proving incomprehensible or include someone's personal information).

5.4: Trade-offs Between Qualities of Documentation Artifacts and Processes

- While each of the qualities in 5.2 and 5.3 is valuable, they cannot all be maximized simultaneously. Maximizing one often comes at the expense of others, necessitating that organizations make trade-offs.
    - Organizations should make trade-offs between qualities of documentation artifacts and processes based on their context and organizational objectives.
- The specific trade-offs can vary, but often include:

- **Informativeness vs. judiciousness:** There are often pieces of information that would be of use to some interested parties, but which other interested parties, possibly including the originating organization, would prefer or even be obliged to protect.
    - E.g., those seeking to assess the suitability of a model that processes personal information might benefit from extensive details about the training dataset, but some of those details could reveal sensitive information about the individuals whose data were included.
        - As they evaluate such trade-offs, organizations should consider user consent regarding data use and sharing, and should respect any prior agreements regarding consent.
    - Similarly, some information can be useful to both legitimate interested parties and malicious attackers, so its release can require careful judgment.
- **Informativeness vs. comprehensibility, maintainability, and keeping processes manageable:** Documentation that is more thorough can provide more of the information interested parties need. At the same time, long or highly detailed documents can be burdensome or overwhelming, potentially making the documentation difficult to understand or use. Such documentation can also tax organizations' resources for creating and maintaining it.
    - This is of particular concern for public documentation, for which non-expert members of the public are frequently part of the intended audience.
    - Organizations can partly address the comprehensibility issue by providing different versions of the documentation that speak to different audiences, at the expense of further burdening the documentation's development and maintenance processes.
- **Informativeness vs. freshness:** The more thorough a documentation artifact is, the more difficult it often is to keep it up to date, at least while keeping the documentation process manageable.
- **Continuous activities vs. keeping processes manageable:** Documenting throughout the life cycle can yield better outcomes, but it can also be more draining for an organization's resources and make documentation workflows more complex.

NIST particularly welcomes input on additional key trade-offs that should be covered, as well as what guidance should be provided for navigating these trade-offs.

Clause 6: AI Dataset and Model Documentation Templates

Clause 6 will provide documentation templates, each of which contains three columns: Column 1 provides numerical identifiers for the fields that make up the AI dataset and model documentation templates; Column 2 provides field names that identify characteristics of AI datasets and models to document; and Column 3 provides a short description of the suggested content of each field.

These templates have been kept at a relatively high level of abstraction, aiming to capture elements that are unambiguously supported by the literature and stakeholder input. Further detail has been deferred to appendices, which offer non-normative examples of how each broad piece of information included as a field in the templates could instead be expanded into multiple granular fields.

6.1: General

- Organizations should use the templates in this clause for publicly documenting AI datasets and models, specializing them as needed.
    - All fields in each template should be provided.
    - The templates seek to achieve the qualities discussed in Clause 5.
        - In particular, they are designed to capture essential information that many interested parties typically need regarding datasets or models.
- The guidance in Clause 5 should be applied when filling out these templates, especially the process guidance.
    - E.g., an organization might fill out templates, perhaps at different levels of detail, at different stages of the lifecycle.
    - The templates naturally cannot capture organizational process elements.
    - Even the artifact guidance might be necessary to determine how best to fill in a field.
- More detail on how each field could be broken apart into more specific fields is provided in Appendices 1 and 2.

NIST particularly welcomes input on the following:

1. The appropriate level of prescriptiveness about use of these templates, i.e., how strongly their use is encouraged.

ISO/IEC language typically distinguishes between purely informative content, guidance ("should" statements), and requirements ("shall" statements), in increasing order of prescriptiveness and amenability to conformity assessment. A recommendation or requirement to use the template could also be softened to say "should consider" or "shall consider."

2. Fields that should be required components for any documentation artifact that claims to instantiate the template.

   In practice, such a requirement would likely be implemented by having a field inherit the same level of prescriptiveness as the use of the overall template—i.e., if use of the template is a requirement ("shall") or a recommendation ("should"), required fields would also be "shall"s or "should"s, respectively.

3. Ways to integrate the proposed dataset and model templates. To ensure consistency, clarity, and ease of use and to reduce redundancy:
   a. The naming conventions, particularly for the field name column, are meant to be as aligned as feasible across the dataset template and the model template.
   b. The dataset and model templates are meant to share field structure and content to the extent feasible, such that relevant data characteristics can easily be kept synchronized between a dataset artifact and a related model artifact.

4. Integration of TEVV results into documentation, particularly fields that would be appropriate to add in the evaluation sections of Appendices 1 and 2. Any fields included would need to be general enough to apply across many kinds of settings, but it might be helpful to have additional suggestions for specific commonly used TEVV points to include.

6.2: Dataset Documentation Template

Identifier	Field Name	Description of Field
1	**Dataset Identifying Descriptors**	Should describe key descriptors that identify the dataset and its core characteristics. For example, organizations might provide information about any subset of the following: ● Formal dataset identifier

Identifier	Field Name	Description of Field
		● Human-readable description ● URL for retrieval ● Dataset and documentation authors and creators ● Contact details ● Version, updates and changes over time ● Digital signature ● Release date ● Date of dataset decommissioning ● Citation ● Attributions
2	**Dataset Intended Use**	Should describe primary applications that the dataset is and is not meant to support. For example, organizations might provide information about any subset of the following: ● Task(s) supported ● Limitations ● Anticipated use ● Out-of-scope use ● Motivation for creation
3	**Usage Rights and Limitations**	Should describe legal terms governing how the dataset may be used, shared, or modified. For example, organizations might provide information about any subset of the following: ● License ● Payments ● Data rightsholders
4	**Dataset Composition and Provenance**	Should describe how and by whom the data has been acquired and processed, including efforts to ascertain or review the quality or nature of the data. For example, organizations might provide information about any subset of the following: ● Data characteristics (e.g., data size, modalities, structure, types, and formats; dataset splits) ● Funding ● PII/confidential data

Identifier	Field Name	Description of Field
		● Data sources ● Sampling procedure ● Consent procedures ● Data flows ● Data enhancement (e.g., synthetic generation of data) ● Pre-processing ● Annotation (e.g., annotators, quality control)
5	**Dataset Evaluation**	Should describe assessment of dataset quality and fitness for use. For example, organizations might provide information about any subset of the following: ● Risk and impact analysis ● Data analysis ● Data validation ● Data visualization
6	**Maintenance and Monitoring**	Should describe how the dataset is updated, shared, and maintained over time. For example, organizations might provide information about any subset of the following: ● Distribution ● Maintainer ● Data storage ● Deprecation plan ● Versioning policy ● Terminology

6.3. AI Model Documentation Template

Identifier	Field Name	Description of Field
1	**Model Identifying Descriptors**	Should describe essential information identifying the model and describing its core characteristics. For

Identifier	Field Name	Description of Field
		example, organizations might provide information about any subset of the following: ● Model or model series identifier ● Model description ● Current model version ID ● Release date ● Point(s) of contact ● Model developer ● Model owner ● Model citation ● Signature ● Model lineage
2	**Model Intended Use**	Should describe the task(s), application domain, capabilities offered, and types of AI system that the model can be integrated into. For example, organizations might provide information about any subset of the following: ● Intended applications ● Intended stakeholders ● Entities using the model ● Capabilities the model offers ● Known limitations ● Out-of-scope use
3	**Usage Rights and Limitations**	Should describe legal terms governing how the dataset may be used, shared, or modified. For example, organizations might provide information about any subset of the following: ● License ● Payment requirements ● Model rightsholders
4	**Model Design**	Should describe how the model is structured and how it is meant to be used within a data processing pipeline from input to output. For example,

Identifier	Field Name	Description of Field
		organizations might provide information about any subset of the following: ● Type(s) of algorithm(s) used ● Model architecture ● Performance metrics targeted ● Input/output formats (e.g., size and modality constraints)
5	**Model Training**	Should describe the training protocol, including hyperparameter configurations. For example, organizations might provide information about any subset of the following: ● Referenced datasets ● Preprocessing steps ● Training data role ● Limitations of training data ● Training protocols
6	**Model Evaluation**	Should provide descriptions of protocol(s) and dataset(s) used to assess the model's quality and fitness for use, analyses of risks and impacts, and the results of evaluation protocols. For example, organizations might provide information about any subset of the following: ● Referenced datasets ● Evaluation data role ● Evaluation of performance limitations ● Quantitative analyses ● Robustness metrics ● Evaluation metrics ● Risks and impacts Organizations might also provide information found in dataset documentation template fields, such as: ● Preprocessing of model/dataset prior to evaluation ● Characteristics of benchmarks

Identifier	Field Name	Description of Field
		● Analysis of quantitative results on benchmarks
7	**Model Maintenance**	Should describe mechanisms to monitor model performance and behavior over time and to report and resolve incidents**.** For example, organizations might provide information about any subset of the following: ● Monitoring mechanisms ● Update mechanisms ● Post-deployment reports ● Incident reporting
8	**Governance**	Should describe frameworks and organizational practices followed in the development of the model, and artifacts resulting from their application. For example, organizations might provide information about any subset of the following: ● Compliance ● Explainability ● Documentation practices and terminology

Appendix 1: Expanded Dataset Documentation Template

Identifier	Field Name	Short Description of Field
1	**Dataset Identifying Descriptors**	**Essential information identifying the dataset and its core characteristics.**
1.1	Dataset Identifier	Provide the name of the dataset, or dynamic datasets bundle, if applicable.
1.2	Dataset Description	Provide a description of the dataset.
1.3	Dataset Authors, Creators, and Developers	Provide name of responsible party: individual(s), entity(ies) creating and/ or developing the dataset.
1.3.1	Dataset Documentation Author(s)	Provide name of entity(ies) preparing dataset documentation artifact.
1.4	Contact Details	Provide information about the point of contact for feedback or questions related to the dataset.
1.4.1	Source URL	Insert direct URL link for the dataset.
1.5	Version	Provide an identifier for the iteration of the dataset or the dynamic dataset bundle being documented, reflecting updates and changes over time.
1.5.1	Digital Signature	Provide a cryptographically signed manifest containing cryptographic hashes for the dataset instances.
1.5.2	Release Date	Provide dataset release date.
1.5.3	Date of Creation	Provide the date of creation of the dataset.
1.5.4	Decommissioning Date	Provide date of out-of-use dataset.

1.6	Citation	Provide technical citation format for referencing the dataset or the dataset bundle.
1.7	Attributions	Provide any notices or attributions required to be associated with the dataset and link to other publicly available documentation about the AI dataset.

2	**Intended Use**	**Motivations for the dataset, including primary applications that it does and does not support.**
2.1	Task(s)	Provide intended application(s) or use case(s) (e.g., classification, translation, summarization).
2.2	Limitations	Describe design constraints, or other limitations that may affect use of the dataset.
2.3	Anticipated Use	Describe expected downstream use, including prior applications or deployments of the dataset.
2.4	Out-of-Scope Use	Provide unintended and prohibited uses of the dataset, including datasets that should not be combined.

3	**Usage Rights and Limitations**	**Legal terms on how the dataset may be used, shared, or modified.**
3.1	License	Insert available license terms governing use of the data. Add any limitations on how the data can be used, modified, or shared, and link to full license text.

		Provide the license in human-readable and machine readable form.
3.2	Payments	Describe whether payment is required to use, modify or share the dataset.
3.3	Data Rightsholders	Identify all parties who have ownership or other interests in the dataset.

4	**Dataset Composition and Provenance**	**Description of how and by whom the data has been acquired and processed, including efforts to ascertain or review the quality or nature of the data**
4.1	Data Characteristics	Provide general information describing data instances.
4.1.1	Data Size	Provide the number of instances included in the dataset (e.g. in tokens or bytes)
4.1.2	Data Modality	Insert type(s) of data included: text, audio, image, video, etc.
4.1.3	Dataset Splits	Describe dataset partitioning and rationale for the split.
4.1.4	Attributes	Describe features or fields present in each data instance.
4.1.5	Data Structure	Describe the underlying structure of the datasets itself: formats, directories, files, etc.
4.1.6	Data Types and Formats	Encoding and representation of attribute values.
4.2	PII/Confidential Data	Describe whether the dataset includes any personally identifiable information or confidential content.

4.3	Data Sources	Provide data acquisition or data generation method.
4.3.1	Data Collection	Identify from where and how the data was obtained, including whether and how the data may be publicly accessible, and if so, on what website or other internet property.
4.3.2	Data Acquisition Date and Method	Identify the date the data was obtained and the method used to obtain it (licensing agreement, publicly accessible download, web scraping, etc.)
4.3.3	Funding	Disclose sources of financial support for the development and maintenance of the dataset.
4.4	Sampling Procedure	Method of selection of data instances.
4.5	Consent Procedures	Provide a description of the process for obtaining informed consent for any collection of human-generated data.
4.6	Data Flows	Provide a description of chain of custody, i.e., how the data went through the system (e.g. source, collection, processing, storage, and the data outflow.)
4.7	Data Enhancement	Provide details about data augmentations to improve or enrich coverage for data; (e.g., stretch, skew the data, any additional features to the data).
4.7.1	Synthetic Generation of Data	Provide a description of the algorithms and software used to generate any synthetic data. Describe how the synthetic data was incorporated into the dataset and what portion of the dataset consists of synthetic content.

4.8	Pre-Processing	Describe initial data transformations before inclusion in the dataset. For example: decisions on data curation processing, data cleaning, filtering, formatting, normalization (handling null values, outliers, etc.)
4.9	Annotation	Describe annotation, or any enrichment performed on the data.
4.9.1	Annotators	Describe the role of annotators and inter-annotator methodology.
4.9.2	Quality Control	Describe methods to ensure consistency and reliability, including to detect and handle errors, mislabeled instances, noise, etc.

5	**Dataset Evaluation**	**Assessment of dataset quality and fitness for use.**
5.1	Risk and Impact Analyses	Provide analyses of risks and impacts of the dataset, including evidence supporting those analyses.
5.2	Data Analysis	Describe analyses to perform data quality distribution, balance, and suitability for the task, e.g., label consistency, data distribution.
5.3	Data Validation	Describe TEVV (testing, evaluation, validation and verification) – e.g., how is it known that the data is valid for the task the dataset is designed and intended for.
5.4	Data Visualization	Provide visual summaries of dataset structure, key properties, and distribution.
5.X	[Other TEVV fields?]	[See call-out box in Clause 6.1]

6	**Maintenance and Monitoring**	**Description of how the dataset is updated, shared, and maintained over time.**
6.1	Distribution	Identify in what way and how the dataset is shared and published, including distribution restrictions.
6.2	Maintainer	Provide the name of the individual or the organization responsible for dataset maintenance, including predecessors of this role.
6.3	Data Storage	Describe how and where the dataset is stored.
6.4	Deprecation Plan	Describe conditions and procedures for phasing out the dataset.
6.5	Versioning Policy	Describe how changes are tracked and documented, including criteria for or frequency of model updates.
6.6	Terminology	Provide technical vocabulary that defines terms of art in a simple and clear manner.

Appendix 2: Expanded Model Documentation Template

Identifier	Field Name	Short Description of Field
1	**Model Details**	**Essential information identifying the model and describing its core characteristics.**
1.1	Model Identifier	Assign a formal identifier to the model or model series. Include the model family, if applicable.
1.2	Model Description	Provide a short description of the model (e.g., in parameters) which could include its size and input and output modalities.
1.3	Model Version ID	Provide a unique identifier associated with the release of the model being documented, with identifiers structured to allow easy determination of whether a version is current.
1.4	Release Date	Provide date of model release.
1.5	Release Methods	Provide the methods of distribution (e.g., list of cloud service providers, product surfaces, websites where weights can be downloaded etc.).
1.6	Point(s) of Contact	Provide feedback mechanism for individuals or organizations who can be contacted for queries, updates, failure modes, flaws, incidents, etc. related to the model.
1.7	Model Developer	Identify entities responsible for developing and maintaining the model.
1.8	Model Owner	Identify individual(s) or organization(s) who have ownership of the model.
1.9	Model Citation	Insert technical citation format for referencing the model, such as a reference to a technical report describing

		the methods used for developing the model.
1.10	Signature	Provide cryptographic signature(s) of the serialized model.
1.11	Model Lineage	Name prior models the model was derived from (e.g., by fine-tuning or distillation) and what mechanisms were used.

2	**Intended Use**	**Description of the task, application domain, capabilities offered, and types of AI system that the model can be integrated into.**
2.1	Entity Using the Model	Identify individuals, groups, organizations deploying the model.
2.2	Intended Applications	Identify primary products, services, or use cases that the model is designed to support.
2.3	Intended Stakeholders	Identify primary individuals, groups, organizations who can benefit from or be impacted by the model and/or its outputs.
2.4	Limitations	Identify exclusions, restrictions on use and/or sharing, as well as constraints that impact performance and/or model reliability. Reference abuse and/or acceptable use policy, if applicable.

3.	**Usage Rights and Limitations**	**Legal terms on how the dataset may be used, shared, or modified.**
3.1	License	Insert available license terms governing use of the model.

		Add any limitations on how the model can be used, modified, or shared, and link to full license text. Provide the license in machine readable form. Indicate with what degree of openness various components of the model (e.g., weights, code, and data) are provided (e.g., proprietary or open access).
3.2	Payments	Describe whether payment is required to use, modify or share the model.
3.3	Model Rightsholders	Identify all parties who have ownership or other interests in the model.

4	**Model Design**	**Description of how the model is used within the data processing pipeline from input to output.**
4.1	Algorithm(s) and Architecture(s) Used	Identify the underlying computational learning paradigm (supervised learning, reinforcement learning, etc.) which help clarify data requirements of the model and the optimization objectives.
4.2	Performance Metrics Targeted	Describe relevant technical design choices and the performance metrics (e.g., precision and recall) that system design and training targeted.

5	**Model Training**	**Description of the training protocol(s) and dataset(s), including hyperparameter configurations.**
5.1-(*)	Referenced Datasets	Provide a reference (Document, URL Link, etc.) pointing to the dataset templates for the dataset used in training. Disclosure

		might also be considered if data is used to train the models for which there is no dataset documentation. This section should be repeated for each referenced dataset used as part of the training process. In the event the model developers made changes to the dataset for training purposes, any overrides or addendums to the information provided in the dataset templates are needed only for those fields in the dataset templates that differ from the base dataset. '(*)' is substituted with a number to identify each reference to a dataset documentation artifact.
5.1.1-(*)	Training Data Preprocessing Steps	See dataset template for training data preprocessing steps, including the order of the training data provided during the training step.
5.1.2-(*)	Training Data Role	Identify purpose of the referenced dataset as used in the training data relative to intended use of the model.
5.1.3-(*)	Limitations of Training Data	Identify quality issues in the referenced dataset that impact model generalizability.
5.2	Training Protocol(s)	Describe the procedures and configurations used to train the model.

6	**Model Evaluation**	**Descriptions of protocol(s) and dataset(s) used to assess the model's quality and fitness for use, analyses of risks and impacts, and the results of evaluation protocols.**

6.1-(*)	Referenced Datasets	Provide a reference (Document, URL Link, etc.) pointing to the dataset documentation artifact for a dataset used in evaluation, if applicable. This section should be repeated for each referenced dataset used as part of the evaluation process. In the event the model developers made changes to the dataset for evaluation purposes, any overrides or addendums to the information provided in the dataset documentation artifacts are needed only for those fields that differ from those in the base dataset's documentation. '(*)' is substituted with a number to identify each dataset documentation artifact. See the following dataset template fields: ● Preprocessing of model/dataset prior to evaluation ● Characteristics of benchmarks ● Analysis of quantitative results on benchmarks.
6.1.1-(*)	Evaluation Data Role	Identify the purpose of evaluation data relative to intended use (e.g., generalization assessment and performance consistency across a variety of use cases). Indicate the test and validation sets, or percentages of data chosen for each.
6.2	Quantitative Analyses	Provide reports of model performance on aggregate and disaggregated performance metrics.
6.3	Evaluation of Performance Limitations	Identify factors that limit the model's performance (e.g., limited dataset, failure modes) and any relevant disclosures about timing and frequency of evaluations.

6.4	Evaluation Metrics	Provide a description of key metrics to assess the model's effectiveness for the intended use, and the model's performance against those metrics, particularly on established benchmarks. Include any metrics that reflect uncertainty estimates, including confidence intervals.
6.4.1	Robustness Metrics	Describe how the model was designed and tested to perform reliably under distributional shifts and adversarial conditions.
6.5	Risk and Impact Analyses	Provide analyses of risks and impacts of the model, including evidence supporting those analyses.
6.X	[Other TEVV fields?]	[See call-out box in Clause 6.1]

7	**Model Maintenance**	**Mechanisms to monitor model performance and behavior over time and to report and resolve incidents.**
7.1	Monitoring Mechanisms	Describe mechanisms the model owner uses to monitor model performance (e.g., analyzing logs, classifiers, feedback). Describe how models are re-evaluated, including to detect drift.
7.2	Update Mechanisms	Describe how models are retrained, including in light of monitoring results.
7.3	Change Log	Provide a link to a list of changes made to the model after deployment.
7.4	Post-Deployment Reports	Describe or reference information about model performance after it is deployed, including major use cases, quantitative analyses, and user feedback.

7.5	Incident Reporting	Add or reference incident reports involving the model, potentially including resolutions of such incidents.

8	**Governance**	**Frameworks and organizational practices followed in the development of the model, and artifacts resulting from their application.**
8.1	Framework Adherence or Compliance	Reference existing compliance and governance frameworks and codes of conduct. Provide links to compliance and governance reports. For example, identify any codes of conduct that have been followed for the AI model (e.g., the G7 Hiroshima AI Process Voluntary Code of Conduct) and any reporting that has been made in connection with any such codes.
8.2	Interpretability	Briefly describe efforts to ensure the model produces its outputs in a way that can be understood and interpreted by humans.
8.3	Documentation	Describe the normative background underlying the documentation artifact, including the organizational practices that yielded the model documentation and how SOPs are documented. Provide technical vocabulary that defines terms of art used in the documentation in a simple and clear manner.

Appendix 3: Example Machine-Readable Model Documentation Template

NIST particularly welcomes input on the potential inclusion of useful examples of a machine-readable documentation template in the standard draft to further substantiate and enhance the wider adoption of dataset and model documentation templates of this standard draft.

Appendix 4: Worked Example of Developing a Dataset Documentation Artifact

NIST particularly welcomes input on any relevant use cases where a dataset and model documentation artifact would be highly valuable for the community, and how to illustrate the use of the AI dataset and model documentation templates in a concrete and cross-domain example.