

AI Documentation Zero Draft Public Comment Template

Commenters on the [initial public draft](#) for the Zero Draft on AI Documentation are encouraged to use this template to provide feedback to NIST.

As noted in the draft document, comments are welcome on any aspect of the draft, particularly the issues raised in the blue boxes. Submissions will become part of the public record and may be subject to public disclosure.

Commenter information

Contact name: Matt O'Shaughnessy	Contact email: moshaughnessy@cdt.org
Organization name (DBA): Center for Democracy & Technology	Preparation or submission date: September 16, 2026

Comments by clause

Instructions: The table below is for comments on specific portions of the draft. For these comments, please indicate in the left column what clause or subclause the comment addresses, and provide the comment, including any suggested changes, in the right column as free text. NIST particularly welcomes specific suggestions for replacement text, where applicable. Create new rows as needed.

Example entries:

Clause/subclause (as specific as applicable)	Comment
Clause 3.5.2	Replace “happy” with “glad.”
Clauses 4.7 and 4.8	Consider combining these clauses. The distinction between them is not strong, and indeed there are edge cases, such as real-time social media video filters, that fall into both categories. The texts of the two clauses could be combined as follows: <Suggested text>

Appendix D	This appendix contains critical information that should be part of the main body. We suggest moving it to become Clause 7.
------------	--

Your comments

Clause/subclause (as specific as applicable)	Comment
Clause 1 (Scope)	Overall, we welcome NIST's efforts to develop this zero draft, which provides a much-needed step toward more standardized guidance on how developers of AI system components should document key properties and potential sources of AI risk. More generally, CDT would like to see AI-related standards developed in venues with as few barriers to broad participation and scrutiny as possible. NIST's "zero drafts" project is a welcome effort to improve the quality and utility of AI-related standards, including by increasing the ability of a variety of stakeholders to contribute.
Clause 1 (Scope)	We are curious about NIST's decision to limit the scope of this document to public-facing AI documentation, rather than also including documentation intended for certain nonpublic audiences. Although some types of model or dataset documentation might be expected to diverge substantially from the considerations described in this document (such as purely internal documentation used to communicate design decisions or requirements among internal teams, where documentation may be more informal, organization-specific, or less comprehensive), many types of non-public documentation (such as dataset or model documentation produced by a vendor for downstream deployers or users) would have considerations that are substantially similar. High-quality documentation for these types of nonpublic audiences is valuable to help actors relying on external AI models and datasets effectively manage risk and advance their own organizational objectives. The guidance and templates provided in this zero draft

	would be helpful for these forms of documentation as well.
Clause 1 (Scope) Clause 2, referenced definition from ISO/IEC 22989 of “AI model”	While we understand that documenting AI systems (rather than just models) would significantly broaden the scope, we remain concerned that this narrow scope excludes the very factors that are most important for consumers of documentation to understand in order to effectively and responsibly integrate AI products into broader systems. AI products are rarely provided as a model object (understood as a model architecture and learned parameters) alone. Rather, AI products (such as user-facing chatbots or API-accessible large language models) typically include pre- or post-processing modules (such as guardrail classifiers), tools, and other components. Standardized documentation that includes these components is critical to allowing downstream users or deployers to make informed adoption, deployment, and risk management decisions. CDT is eager to see future work from NIST that moves toward system-level documentation guidance. NIST notes that concepts “beyond the model object” have been excluded from this document “given the dearth of literature for standardization to build on” (“Note for Reviewers,” line 165). More public information or analysis from NIST describing specific gaps in pre-standardization research or questions that future literature should address to best enable efforts to create system-level documentation standards would be helpful. In determining the appropriate scope of the document and term “AI model,” we suggest targeting the typical scope of AI models provided by AI developers (for example, through API access) to downstream deployers, understanding that this is a broad category that does not define a bright line definition of “AI model.” NIST notes that it “has not identified a principled basis for scoping in certain system components, such as those sometimes referred to as part of a model, while scoping out others” (“Note for

	Reviewers,” line 165). However, we note that “model” as defined in ISO/IEC 22989:2002 (Clause 3.1.23) (“physical, mathematical or otherwise logical representation of a system, entity, phenomenon, process or data”) itself appears to be ambiguous on the question of what parts of an “AI system” lie within or outside of the scope of the “AI model”; it may be necessary to accept a degree of ambiguity in the scope of “AI model” and the model documentation template. The “conditionally recommended” designation in the AI model documentation profile (Appendix A) provides an ideal mechanism to include fields that may not be applicable to all models, mitigating the adverse impact including more fields that may not be universally applicable across models.
Clause 3 (Dimensions of potential outcomes of public-facing documentation processes)	Generally, we are concerned that this section (and the draft) is framed too narrowly around advancing the organizational objectives of the provider. While we appreciate the inclusion of aspects generally related to the objectives of end-users, other external parties, or the public through considerations such as “accountability,” “privacy preservation,” and “trust” (Clause 3.4), describing these important elements only as considerations that could indirectly advance the provider’s organizational objective appears to open the possibility for providers to point to this eventual standard to help justify withholding valuable information from documentation even when it comes at the cost of the public interest.
Clause 3 (Dimensions of potential outcomes of public-facing documentation processes)	We suggest adding more explicit discussion of how robust documentation artifacts can improve the ability of downstream actors adopting or deploying AI datasets to more effectively manage risks in their own activities.
Clause 4.2.5 (Artifact freshness)	We support the inclusion of the consideration that “Older versions of documentation might remain useful[.]” This consideration might benefit from a short description of reasons older versions of documentation might be useful. For example, older versions of documentation allow downstream developers, users, or the public to track how aspects of a particular

	dataset or model has changed over time, or support downstream developers and users who continue to use legacy datasets or models.
Clause 4.2.6 (Artifact interoperability)	We appreciate the inclusion of the note that organizations should provide both machine-readable and human-readable documentation artifacts. Machine-readable documentation artifacts support valuable independent research. We suggest specifying in this section that organizations should maintain consistent (or at least interoperable) structure across documentation artifacts for different versions of the same dataset or model. Where exact consistency is not possible or feasible, organizations should make an effort to improve clarity in how the structure and content has changed across versions. We also suggest specifying in this section that organizations should make an effort to improve interoperability in metrics, benchmarks, evaluations, and other measures. For example, using similar (or comparable) metrics allows more accurate comparison of how various dimensions of downstream performance might be affected by different versions of a dataset or model, or by different datasets or models altogether. This should also include a specification that organizations should clearly describe challenges to interpretability or interoperability, or limitations to comparing various dimensions of performance across different datasets or models.
Clause 4.2.7 (Artifact findability and availability)	We suggest including a brief mention that previous versions of documentation artifacts (and documentation artifacts for previous versions of datasets or models) should be made available with minimal barriers to access.
Clause 4.3.3 (Provide organizational support)	We support the inclusion of establishing institutional incentives for prioritizing documentation, and suggest adding incentives related to organizational subcomponents (for example, including aspects related to documentation in KPIs or other measures of

	team performance) in addition to individual-focused incentives.
Clause 5.1 (Documentation templates – General)	Line 610 reads “The guidance in Clause 4 should be followed when using these templates.” The following lines in this paragraph similarly read that guidance in subclauses under Clause 4 “should” be followed by organizations. We suggest modifying these statements to “shall.” The guidance in Clause 4 already offers extensive flexibility for organizations to make trade-offs among competing factors to advance organizational objectives.
Clause 5.1 (Documentation templates – General) – Note to reviewers at line 618	NIST inquires whether and how it should provide “concrete examples” for templates (“Note to Reviewers,” line 618). On the one hand, we believe that filled-in template(s) for representative dataset(s) and model(s) would be helpful to promote understanding of potentially relevant content for each field. However, we are concerned that example filled-in templates that offer only a limited amount of detail could contribute to a norm or expectation for organizations to limit the amount of information provided in publicly-available documentation artifacts. As a result, we suggest that should NIST provide example filled-in templates, at least one example should have significant detail, appropriate for a high-use, high-complexity, or high-impact dataset or model. (Additional less detailed example(s) could be tailored for datasets or models where a more limited level of detail is appropriate.) Framing text that makes clear that example templates provided by NIST should not be used to justify minimizing the extent of publicly-available documentation could also be helpful.
Clause 5.2 (Dataset documentation template) – Field ID 2 (Intended use) Clause 5.3 (Model documentation template) – Field ID 2 (Intended use)	We support the specification that this field should potentially include a description of “applications the dataset [model] is known to be unsuitable for.” It is particularly important for downstream developers or users to be able to identify when any particular application is not suitable for a dataset (or model). We suggest expanding this language to more explicitly recommend including a description of factors or

	general traits of applications that the dataset (or model) might be unsuitable for.
Clause 5.2 (Dataset documentation template) Clause 5.3 (Model documentation template)	We suggest NIST adopt a lower threshold for marking fields as “Recommended” rather than “Optional,” or provide more detail in general notes clarifying that datasets (or models) that are likely to be more widely used or lead to greater downstream impacts should specify more detail. Specifying that there should be “no particular expectation that the field will be populated” sets an incredibly low bar for important information including intended use, evaluation information, or usage rights. This risks contributing to norms or expectations of publicly-available documentation that is too terse to be useful.
Clause 5.2 (Dataset documentation template) – Field ID 5 (Evaluation) Clause 5.3 (Model documentation template) – Field ID 6 (Evaluation)	We suggest describing the benefits or importance of including structured data about evaluation protocols and datasets that would allow potential users to make comparisons across datasets or models.

Document-wide comments

Instructions: This section is for comments that pertain to the entire document, not specific clauses or subclauses. Create new rows as needed.
