

Purpose. Document an AI model or automated decision system so its behavior and decision logic are transparent, auditable, and understandable to nonexperts.

What it captures. The intended use and out-of-scope uses, performance against the metrics that matter, the go-live criteria, known limitations, ethical considerations and safeguards, the human-oversight path, and the owner and review.

When to use it. Before an AI model or automated decision system goes live, and whenever its rules, data, intended use, or safeguards change.

Common mistake. Documenting accuracy while omitting the limitations and the human-oversight path, which is the part a nontechnical reviewer most needs. A model card without its failure modes builds false confidence rather than trust.

Decision rule. A model or automated decision system goes live only with a documented human-oversight path and at least one named limitation. If no one can explain how the system fails or how a person can override it, it is not ready for the floor.

Model or System Name**Version and Date****Intended Use** *what it assists with, and whose judgment it does not replace***Out-of-Scope Uses****Performance** *against the metrics that matter, with the evaluation set described***Go-Live Criteria** *the error-cost, edge-case, and harm-and-fairness thresholds the system must meet before it goes live***Known Limitations** *the failure modes a nontechnical reviewer most needs to know*

-
-
-

Ethical Considerations and Safeguards**Human-Oversight Path** *how a person can override the system, and how overrides are monitored***Owner****Review Trigger**