

Appendix I — AI Safety and Learning Log

Joe Gleason · Advance review draft 3.7 — October 2026

Respond to problems and preserve useful improvements

Use with: Chapters 9, 10, 13, and 14.

This log follows the source specification for questionable output, consequences, correction, and resulting change. It does not replace the organization's immediate care, safety, or incident-response procedures. Attend to the person first; record and review the learning through the appropriate route.

I1. One issue or learning entry

Field	What to record
Event and version	Date, task, workflow/model configuration, source set, and output version.
How it was noticed	The question or observation raised by a patient, participant, caregiver, volunteer, staff member, reviewer, or other source.
What happened	The exact problem or useful observation, with relevant material stored only in its approved location.
Exposure and consequence	Was it reviewed, shared, or acted on? Who may be affected? Distinguish observed consequences from possible ones.
Immediate response	Correction, appropriate contact, containment, or a decision to pause—by whom and when.
Cause and uncertainty	What is established, what is suspected, and what still needs investigation.
Decision and change	The named person's decision and the actual revision to sources, instructions, workflow, training, or access.

Field	What to record
Verification	How the change was tested, what remains unresolved, and the next review date.
Credit and permitted learning	How contributors wish to be acknowledged and what lesson can be shared without exposing private information.

Record problems that were caught before publication as well as those reaching users. Do not reward only positive reports or treat a person raising a concern as the problem.

12. Optional live and post-encounter assistance

Agree to live assistance before use: what may be captured, who receives it, how to pause, and the human-led alternative. The facilitator checks prompts about unresolved questions; the assistant does not infer emotions or decide who is right.

Afterward, organize only permitted feedback, preserve different accounts, and record the change people actually made—not a claim that the AI learned.

Record the decision, actual change, and remeasurement, including difficult results. Keep this log outside the software when needed; Appendix J records the dated capability limits.

Decision this tool supports: Has the risk been addressed, has the correction worked, and what should the organization learn or share?