A ONE-PAGE GUIDE FOR DOCTORAL RESEARCHERS

The *research harness*

A structured specification of an AI agent's operating context in doctoral research.

THE PROBLEM. AI agents now contribute to doctoral work, but mostly under ad hoc arrangements: individual exchanges with the agent, governed by no shared account of what it is doing across the project as a whole. The familiar troubles follow from that absence — work running ahead of the thinking it depends on, the quiet offloading of interpretation, a project drifting as each plausible suggestion accumulates, an agent that tends to agree with you, inconsistent answers to the same question, and, later, no way to say who contributed what. Institutional policy and AI-literacy training do not reach this, because the gap sits between the institutional rule and the individual exchange.

THE IDEA. A harness is a structured specification of the agent's operating context, borrowed from software engineering, where agents working without one produced the same kinds of trouble. It has seven components. You can enter it **thin** — a single sentence under each — and grow it as the work demands; each extension is a small piece of supervised doctoral work. Throughout, you stay in the loop: the agent contributes to the inquiry, but you remain its analyst, its judge, and its author.

The seven components

Build roughly in this order; revise iteratively across the project.

1 Knowledge base

The curated material the agent may work within — your question, framework, methodology, literature, permitted data, ethics approval. If it isn't here, it doesn't exist for the agent.

Start · Point the agent at documents you already have, and state explicitly what it must never see (sensitive or identifiable data).

2 Interpretive permissions

The analytic rules for how the agent reasons from that material — what counts as a legitimate inference and what is overreach in your tradition.

Start · Name your tradition, add two or three "must not" rules you can commit to now, and require the agent to offer options rather than recommendations.

3 Tools

The capabilities the agent may invoke — literature search, reading transcripts, interrogating a dataset, retrieval from a reference manager. Named as capabilities, not products.

Start · List what it can do and, just as deliberately, what it cannot; revisit when you change platforms.

4 Authority

What the agent may do with those capabilities, in three tiers: *autonomous* (reversible, inspectable), *supervised* (it proposes, you judge — most work lives here), and *reserved* (never delegated, e.g. final interpretation).

Start · Sort concrete actions into the three tiers with your supervisor.

5 Scope register

A designated place to park interesting-but-off-topic material the agent surfaces, so it is preserved rather than chased or lost.

Start · Tell the agent to log such items to one file; review them when the project's direction shifts.

6 Process record

A versioned log of what was asked, produced, and decided — and why. It is the agent's external memory and a discipline against offloading.

Start · At each session's end have the agent draft the update; you edit it and supply the rationale yourself.

7 Amendment protocol

How the harness changes. It separates *exceptions* (one-off steps outside the harness, logged and addressed) from *amendments* (deliberate changes to the harness itself).

Start · When something shifts, decide which it is and record it; discuss substantive amendments with your supervisor.

→ Putting it to work

Load the harness into your AI tool's project workspace at the start of a session; bring only the components a given task needs. You brief the agent on it — it is not a setting of the system.

Remember · Read against your supervisor's response — many components only become meaningful that way.

THE FORM A folder of plain **markdown files**, named by component and versioned with Git or simply dated copies, kept alongside your project materials. It is a research artefact first and a technical one second — no software skills required, nothing you couldn't pick up in an afternoon. Begin with one sentence per component, build it alongside the work, and let it mature as the project meets cases the first version did not anticipate.

CITE AS Rowe, M. (2026). *The research harness: a framework for bounded AI use in doctoral work.* Open Science Framework. https://doi.org/10.35542/osf.io/mwhgz_v1