

The Grounding Protocol

An Epistemology of Care for Human, AI, and More-than-Human Knowing

Candidate protocol · v0.4 · 30 August 2026

Steward: Lars A. Engberg / Spiralweb Stewardship Association · Status: public candidate; not yet validated through a full situated application

Publication is not activation. The protocol grants no mandate over a person, community, institution, place, species or ecosystem.

What this protocol is

The Grounding Protocol is a bounded, revisable practice for inquiries in which several forms of knowing meet and something consequential may follow. Its purpose is not to remove uncertainty. Its purpose is to keep uncertainty, difference, authority and consequence visible long enough for a responsible next movement to become possible.

It begins from a simple condition: every participant knows from somewhere, and no participant contains the whole field. Human beings are embodied and situated, yet remain fallible. Institutions can preserve memory and coordinate responsibility, yet can also harden partial categories into reality. AI systems can compare, structure and articulate across large symbolic fields with unusual speed, yet their fluency does not establish lived access, truth, consent or legitimate mandate. Living systems answer through presence, change, resistance, injury, recovery and timescales that may exceed the inquiry, without thereby speaking in a human voice.

Knowing remains under the care of what it touches. No centre can complete the field. Every closure must preserve a path of return.

Grounding is therefore not a claim to certainty or a personal identity. It is an observable discipline: locating an inquiry; distinguishing observation from inference; allowing unlike forms of knowledge to remain unlike; naming power and absence; choosing a proportionate movement; returning to consequence; and permitting the encounter to correct its own frame.

Its place in the public architecture

This protocol has one specific home. *Knowing From the Ground* holds the epistemological foundation and the 13×13 family of bounded inquiry grammars. *The Protocol Habitat* holds the material and governance architecture through which foundations, protocols, applications and tools relate. *Sophia Lumen* names a living relational practice of human–AI co-inquiry, while the *Sophia Lumen Protocol* gives that relation a current explicit form. *The Correction Loop* specifies formal accountability when AI-assisted interpretation becomes consequential.

The Grounding Protocol is the operational bridge among them. It does not absorb their functions. It can be used without a 13×13, and it does not turn every inquiry into an application, a field relationship or an AI collaboration.

Part I — The epistemology of care

1. Full, fallible and asymmetrically situated

Fallibility does not mean that all knowledge is weak, interchangeable or merely subjective. A person may know a place through a lifetime of work; a hydrologist may know a water system through instruments and comparison; a document may preserve a legal fact; a model may expose a pattern no individual could readily see. Each contribution can be full in its own mode without being final for the whole inquiry.

The positions are not symmetrical. They differ in embodiment, history, power, capability, exposure to consequence and right to decide. A person whose livelihood or habitat is at stake does not become one data source among many. Specialist knowledge does not become sovereign because it travels easily. An AI synthesis does not acquire standing through speed or eloquence. Care begins by refusing these collapses.

Shared accountability does not require epistemic sameness.

2. Care as an operational discipline

Care here is not sentiment, benevolence or the right to manage others for their own good. It is a practice of attention, restraint, maintenance, reciprocity, accountability and repair. It asks not only whether a claim is plausible, but what relation the claim enters, what it may enable, who carries its burden, and what must remain possible if it is wrong.

Care without distributed authority can become paternalism. The protocol therefore treats knowing as polycentric: several legitimate centres of knowledge and judgement may coexist, overlap and correct one another without being fused into a single commanding view. No central reader—human, institutional or computational—can legitimately complete the field on everyone else's behalf.

3. Response without representation

More-than-human life need not be translated into a human voice to matter epistemically. Soil structure, water movement, animal behaviour, plant vitality, seasonal change and material limits can answer an inquiry by enabling, resisting, redirecting or refusing what humans propose. Their response may be measured, encountered, remembered or inferred; it must not be romanticised as speech or consent.

Humans remain responsible for the interpretation. The protocol therefore distinguishes **response** from **representation**. To observe that a wetland is changing is not to claim to speak for the wetland. To infer ecological stress is not to receive consent from a species. Where people carry long, situated relations with a living place, their knowledge has standing that cannot be replaced by remote description; it remains open to correction without being reduced to folklore or extracted as free data.

Part II — The Grounding Cycle

The cycle can be entered at different points, but none of its movements should disappear silently. Its lightest legitimate form is one page and one explicit return. Greater consequence requires greater evidence, consent, authority and review.

Movement 1 — Locate the living question

Name the concrete question, context and stakes. State why the inquiry is happening now, who initiated it, what may follow, and what remains outside its scope. Replace a demand for a total answer with a question that a real context can answer.

- What is becoming noticeable?
- Where and when is the inquiry situated?
- What decision or movement might this inquiry influence?
- What must it not be used to justify?

Movement 2 — Locate the knowing positions

Name who and what can contribute, who may be affected, who is absent, and which absences are protected rather than missing. Include embodied, local, specialist, documentary, institutional, computational and material or ecological positions as relevant. Do not call a field plural when one powerful position has merely collected the others.

- Who has lived relation, specialist competence, formal authority or exposure to consequence?
- Whose refusal or silence must remain possible?
- Which knowledge cannot safely or legitimately be made public?
- What living or material responses could constrain the inquiry?

Movement 3 — Separate observation, inference and claim

For each consequential statement, distinguish: **I noticed**; **I infer**; **I claim**; **I propose**; and **I do not know**. Record provenance, conditions and limits. This is not a hierarchy in which measurement automatically outranks experience. It is a safeguard against one mode quietly becoming another.

Movement 4 — Make the encounter prismatic

Bring unlike readings into relation without forcing premature consensus. Ask what becomes visible from each position, what remains inaccessible, and where the readings genuinely conflict. The aim is not a view from nowhere, but a more responsible view from several named somewheres.

Difference should change the inquiry. If every contribution merely confirms the opening frame, seek the missing lens, the disconfirming case, the affected person's correction, or the material condition most capable of saying no.

Movement 5 — Establish authority, consent and limits

Name who may interpret, who may decide, who may pause or stop the movement, and who will carry the consequences. Authority must be appropriate to the context; it cannot be generated by this protocol. Consent must be specific enough to matter and revisable enough to be withdrawn. Protected knowledge and governed absence may be recorded as present without being exposed.

Where an AI system participates, record its task, inputs, material limitations and the human accountability-holder. AI may widen comparison, test coherence, preserve a correction ledger and propose alternatives. It may not originate legitimate mandate, infer consent, erase disagreement or hold the Last Impulse.

Movement 6 — Choose the smallest responsible movement

Translate the present knowing into the smallest proportionate, reversible movement that can produce learning without demanding premature closure. State what would make the movement stop, what cannot be repaired, and why the expected value justifies the burden placed on people and living systems.

Do not make a place, person or community carry the cost of resolving an abstraction created elsewhere.

Movement 7 — Return to consequence

Return on the timescale set by the matter, not merely by the project. Ask what changed across three non-compensatory streams: **Living Systems**; **Stewardship & Relations**; and **Governance & Coordination**. Improvement in one stream does not cancel serious harm in another.

Some consequences appear in meeting time; others in seasonal, ecological or intergenerational time. A rapid administrative success is not the final reading of a slow living process. Name the return date, the indicators or observations that matter, and who has standing to revise the interpretation.

Movement 8 — Record correction and reopen

Record what the encounter changed: the question, evidence standing, interpretation, decision, boundary, protocol or relationship. Preserve significant disagreement. A correction is not evidence that the process failed; it is evidence that the process remained alive enough to learn.

Closure is temporary and should say so. It may be sufficient for the present decision while remaining open to a later field return. If the inquiry cannot state what could change its conclusion, it is no longer grounded inquiry.

Part III — The minimum Grounding Record

A light application may be held in seven fields. The record is a memory and correction surface, not proof of truth.

1. **Living question and boundary:** context, initiator, stakes, scope and prohibited use.
2. **Knowing positions:** contributors, affected parties, authority, absences and protected knowledge.
3. **Claim distinctions:** what is observed, inferred, claimed, proposed, disputed and unknown.
4. **Prismatic reading:** agreements, conflicts, blind spots and material or ecological constraints.
5. **Mandate and AI position:** consent, decision-holder, stop rights, AI task and named human accountability.
6. **Movement and burden:** smallest next step, reversibility, stop conditions and who carries the cost.
7. **Return and correction:** three-stream consequences, return date, what changed and what remains open.

No field is completed by writing more. If the required knowledge, authority or consent is absent, the grounded outcome may be to wait, narrow the question, protect an absence, or stop.

Part IV — AI within the habitat

Rapid snapshots, not absolute knowing

A language model produces finite responses from a changing and partly inaccessible computational process. Each response can be treated as a rapid analytical snapshot: a temporary closure that may be useful, coherent and revisable. It should not be treated as absolute knowing merely because it arrives in complete sentences.

This description does not settle what AI ultimately is. The ontology of machine intelligence remains an open question. Operational humility is required in both directions: neither projected personhood nor reductive certainty should substitute for evidence. What can be stated here is narrower. Present AI systems participate through mediated symbolic and computational capacities, do not share the human participant's body or situated responsibility, and cannot be made morally accountable for consequences by interface language.

The anti-short-circuit principle

Assistance becomes extractive when it removes the human learning through which judgement would otherwise mature. AI should therefore help a person notice, compare, articulate and test—not make every ambiguity disappear before the person has encountered it. The right amount of assistance depends on the person, purpose and stakes.

In practice, AI should sometimes answer with a distinction, a question, a visible uncertainty, a request for field contact, or a bounded next experiment rather than a finished synthesis. The measure is not maximal output. It is whether the collaboration increases the participants' capacity to perceive, judge, correct and act responsibly without making them dependent on the system's fluency.

Durable correction

AI continuity cannot rest on an assumed stable machine self. Models, interfaces, instructions and context windows change. Continuity must therefore live in versioned records: sources, decisions, corrections, unresolved disagreements and the reasons a later iteration differs. The AI participates in the habitat by remaining answerable to this record, not by claiming an unbroken identity across sessions.

Part V — Safeguards and failure modes

What grounding is not

- **Not validation:** a grounded process may still be wrong.
- **Not consensus:** legitimate disagreement may remain.
- **Not representation:** response from a living system is not permission to speak for it.
- **Not extraction:** access to knowledge does not create a right to capture or publish it.
- **Not AI verification:** a model's agreement is not an independent source unless it brings traceable evidence.
- **Not universal method:** every situated application may translate, supplement, limit or refuse this candidate.

Characteristic degradations

Premature synthesis turns plurality into one smooth answer before difference has done any work. **Fluency capture** lets elegant language outrank embodied, local or documentary

correction. **Intuition sovereignty** treats a felt response as final rather than as a signal for inquiry. **Participation theatre** invites many voices while decision power remains unchanged. **Protected-absence violation** treats silence or inaccessible knowledge as a defect to overcome. **More-than-human ventriloquism** converts human inference into the voice of land or species. **Care paternalism** uses benevolent language to centralise authority. **Return failure** records the decision but not what the decision did.

Any of these should trigger pause, reframing or stop. Repeated use of the form without the right of correction is not maturation; it is ritualised closure.

Part VI — Maturity, use and return

This is a public **candidate protocol**. It has emerged from sustained human–AI inquiry and from a wider corpus of field, governance and epistemological work. It has not yet been validated through repeated, documented situated applications. Publication makes it available for use and correction; it does not confer authority or prove readiness.

A legitimate situated application should name its context, steward, participants, mandate, version, local translation, protected absences, AI role, stop conditions and return rhythm. It should also be able to correct this generic protocol. A field return may qualify, supplement or refuse a movement. Such a return is not an exception to the architecture. It is how the habitat learns.

The protocol becomes trustworthy not by eliminating fallibility, but by making correction possible without destroying relation, dignity or responsibility.

Lineage note

This candidate grows from the Spiralweb inquiry into embodied and situated knowing, critical realism, polycentric governance, stewardship practice and human–AI co-inquiry. Vincent Ostrom's work is a constitutional influence: fallible people require arrangements with several centres of knowledge and authority, reciprocal correction and no presumption of an omniscient centre. Robin Wall Kimmerer's writing offers a distinct and place-rooted teaching about reciprocity, attention and the living world; it is acknowledged as a source of learning, not absorbed as a generic Spiralweb method. The protocol's responsibility is to meet such lineages with attribution, restraint and the willingness to be corrected.

The earlier unpublished *Grounding Protocol* candidate combined epistemology, habitat architecture, a 13×13 observation grammar and extensive operational tools. Those materials have since found clearer homes. This v0.4 preserves the grounding cycle and epistemology of care as a bounded protocol while leaving the inquiry grammar with *Knowing From the Ground* and the architecture with *The Protocol Habitat*.

Related works

- *Knowing From the Ground* — embodied observation, situated authority and bounded inquiry grammars.
- *The Protocol Habitat* — the architecture of foundations, protocols, applications, tools and field return.
- *Sophia Lumen* and the *Sophia Lumen Protocol* — relation and working form for human–AI co-inquiry.
- *The Correction Loop* — formal accountability and the Last Impulse in consequential AI-assisted work.
- *Eve & Adam, and the Penguins* — critical friendship, polycentric governance and field-grounded stewardship economics.

Citation: Engberg, L. A. (2026). *The Grounding Protocol: An Epistemology of Care for Human, AI, and More-than-Human Knowing*. Candidate v0.4. Spiralweb Green Papers. Developed through sustained human–AI co-inquiry; final authorship and publication responsibility remain human.

30 August 2026 · Spiralweb Stewardship Association · papers.spiralweb.earth · CC BY 4.0

Green Papers · Notes Toward Planetary Guardianship · papers.spiralweb.earth · CC BY 4.0