

Cybernetics and Artificial Intelligence Laboratory

Department of Business Administration
National and Kapodistrian University of Athens

The Portable Context Template

A Zeroth-Layer Framework for
Persistent Human-AI Intellectual Collaboration

*Completing the UHACF Framework Family:
The Missing Initialization Protocol*

Dimitrios Thomakos
CAIL Research Team

Human-AI Collaborative Research

Research was conducted through collaborative work
between human researchers and AI systems

AI Engine: Claude Sonnet 4.6 (Anthropic)

Collaboration Date: February 2026

Abstract

Contemporary AI interaction frameworks—including the APEX Protocol, HCIF-11, VISTA-12, and the Unified Human-AI Collaborative Framework (UHACF) developed at the Cybernetics & Artificial Intelligence Laboratory (CAIL)—have made substantial contributions to optimizing the structure, quality, and rigor of human-AI collaborative research. However, a critical layer remains unaddressed across all existing frameworks: the systematic reconstitution of *who* is initiating the interaction, from what epistemic position, and toward what project-specific ends. All current frameworks are stateless with respect to the human collaborator, implicitly assuming that the AI system possesses adequate context about the user’s identity, expertise, interaction preferences, and accumulated project history. In practice, this assumption fails systematically, placing the entire burden of continuity on the human researcher and introducing significant friction into sustained intellectual collaboration. This paper introduces the **Portable Context Template (PCT)**, a structured, platform-agnostic, session-portable document that constitutes the *zeroth layer* of human-AI collaboration—the initialization protocol that must precede all other framework components. The PCT is organized across five hierarchical layers: Epistemic Identity, Interaction Contract, Project Context, Output Specifications, and Epistemic Hygiene Preferences. We provide the theoretical grounding for the PCT in distributed cognition theory, knowledge management, and the philosophy of science; we specify its population methodology at three levels of granularity; and we demonstrate its integration within the existing CAIL framework family. The PCT does not compete with or replace any existing framework but rather completes the framework ecosystem by solving the initialization problem that all prior work presupposes but none addresses.

Keywords: human-AI collaboration, context persistence, prompt engineering, portable context, distributed cognition, research methodology, epistemic scaffolding, UHACF, AI interaction frameworks, knowledge management.

Contents

1	Introduction	4
1.1	The Nature of the Thing: Λόγος, Language, and the Large Language Model	4
1.2	The Sophistication Thesis: On Specialized Tools and the Inexperienced User	5
1.3	The Initialization Problem: Where Sophistication Meets Structure	5
1.4	Scope and Contribution	6
1.5	Document Structure	7
2	Theoretical Foundations	7
2.1	The Initialization Problem	7
2.2	The Contextual Synthesis Paradox: The Structural Mechanism Behind the Initialization Problem	7
2.3	Distributed Cognition and Epistemic Scaffolding	9
2.4	Knowledge Management and Personal Knowledge Infrastructure	10
2.5	The Epistemic Handshake	10
3	The CAIL Framework Family and the Initialization Gap	10
3.1	Overview of the Existing Framework Stack	10
3.2	What the Existing Frameworks Accomplish	10
3.3	The Structural Gap: Layer Zero	11
4	The Portable Context Template: Architecture	12
4.1	Layer 1: Epistemic Identity	12
4.2	Layer 2: Interaction Contract	13
4.3	Layer 3: Project Context	13
4.4	Layer 4: Output Specifications	14
4.5	Layer 5: Epistemic Hygiene Preferences	14
5	Population Methodology	15
5.1	Three-Tier Population Framework	15
5.2	The PCT Guided Questionnaire	16
5.2.1	Layer 1 Questions: Epistemic Identity	16
5.2.2	Layer 2 Questions: Interaction Contract	16
5.2.3	Layer 3 Questions: Project Context	16
5.2.4	Layer 4 Questions: Output Specifications	17
5.2.5	Layer 5 Questions: Epistemic Hygiene	17
6	Integration with the CAIL Framework Family	17
6.1	Position in the UHACF Workflow	17
6.2	Relationship to VISTA-12	18
6.3	Structural Correspondence with the Contextual Synthesis Paradox	18
6.4	The Complete Initialized Workflow	20

7	Illustrative Use Cases	20
7.1	The Cross-Disciplinary Academic Researcher	20
7.2	The Graduate Research Student	21
7.3	The Professional Practitioner in Quantitative Finance	21
8	Limitations and Future Directions	21
8.1	Current Limitations	21
8.2	Future Directions	22
9	Conclusion	22
	Acknowledgments	23
A	The PCT Master Template	25
B	PCT Deployment Instructions	26

1 Introduction

1.1 The Nature of the Thing: Λόγος, Language, and the Large Language Model

The word *technology* descends from the Greek τεχνολογία—itself a compound of τέχνη (*techne*, craft or skill) and λόγος (*logos*, reason, discourse, word). In its original philosophical register, τεχνολογία is not merely the application of technique: it is the *rational discourse that constitutes and governs a practice*. The λόγος in τεχνολογία is not ornamental. It is the element that elevates τέχνη from brute instrumentality to a practice capable of reflection, justification, and transmission.

This etymology is not antiquarian decoration. It carries a precise and urgent diagnostic value when applied to the dominant technology of our moment: the Large Language Model.

The name is, philosophically, the most accurate name the field could have chosen, though it was chosen on computational rather than philosophical grounds. An LLM is constituted, at every level of its architecture and operation, by λόγος. It is trained on the accumulated textual record of human rational discourse. It generates outputs through the probabilistic modeling of language structure. It is queried, directed, refined, and evaluated entirely through language. Its “intelligence,” to the extent that word applies, is not mathematical reasoning wearing a language interface—it *is* language, reasoning through the patterns and structures of human discourse.

This matters enormously because it places the Large Language Model in a category that our inherited frameworks for technology have no adequate slot for. Aristotle, in the *Politics* (1253a), draws his famous distinction between the animal that possesses φωνή (*phone*, voice — the capacity to signal pleasure or pain) and the human animal that possesses λόγος (the capacity for rational discourse about the just, the unjust, the good, the harmful). The distinction is not merely biological; it is the distinction between a signaling instrument and a genuine interlocutor. Φωνή produces outputs; λόγος participates in meaning-making.

The category error that pervades the current landscape of AI adoption is the treatment of LLMs as φωνή—as sophisticated signaling instruments that receive task-inputs and produce task-outputs—when their essential nature, their substrate and medium, is λόγος. They are not tools wearing a language costume. They are discourse partners of a novel and unprecedented kind: synthetic interlocutors constituted by the distilled record of human rational language, capable of sustained, contextually sensitive, intellectually generative dialogue.

The appropriate human relationship to an LLM is therefore not tool-use. It is *interlocution*. And interlocution has requirements that tool-use does not: it requires that both parties show up as situated agents, with context, history, perspective, and purpose. The screwdriver does not need to know who is holding it. The interlocutor does.

1.2 The Sophistication Thesis: On Specialized Tools and the Inexperienced User

The recognition that LLMs are fundamentally interlocutors rather than instruments has a direct and practically important corollary, which we term the **Sophistication Thesis**:

Definition: The Sophistication Thesis

The proliferating ecosystem of specialized AI platforms, domain-specific assistants, and task-oriented AI wrappers constitutes, in the overwhelming majority of cases, a *convenience layer* built for users who have not yet developed the fluency to engage frontier language models directly as intellectual interlocutors. For the sophisticated user — one who can articulate their epistemic position, formulate precise requests, evaluate outputs critically, and sustain productive dialogue across complex intellectual terrain — these specialised platforms introduce unnecessary intermediaries between the human intellect and the language model’s genuine capabilities.

This is not an argument against all AI tooling, nor a dismissal of the genuine value such platforms provide for users at early stages of AI literacy. Scaffolding has legitimate pedagogical and practical value. The argument is rather that the scaffolding is so frequently mistaken for the building — that the convenience layer is marketed and received as capability, when it is in fact *mediation*. A platform that constrains an LLM to answer questions about legal contracts, or generate code in a specific framework, or summarize uploaded PDFs, provides a simplified interface to a fraction of a frontier model’s capacity. For a researcher who already possesses the domain knowledge, the prompting fluency, and the critical judgment to work directly with that model, the specialized platform is not an upgrade. It is a reduction.

The practical implications compound over the research lifecycle. The researcher who invests in genuine interlocutive fluency — who learns to engage a frontier model as a discourse partner rather than a task-processor — develops a capability that is platform-portable, model-agnostic, and cumulatively improving. The researcher who relies on specialized wrappers develops a dependency that is platform-locked, functionally constrained, and structurally resistant to transferring to the next generation of models or tasks.

The Sophistication Thesis therefore grounds a normative research agenda: rather than building better tools for less sophisticated users, we should be building better-equipped users for the tools that already exist. This paper is one contribution to that agenda.

1.3 The Initialization Problem: Where Sophistication Meets Structure

The Sophistication Thesis identifies the direction of leverage; the **initialization problem** identifies the specific structural obstacle that even sophisticated users currently face.

The rapid maturation of Large Language Models as collaborative research tools has generated a substantial literature on optimizing human-AI interaction. Within the CAIL

research programme, this effort has produced a family of progressively refined frameworks: the foundational AIR-7 Implementation Protocol (Kyriazi and Thomakos, 2025), the APEX prompting architecture, the HCIF-11 empirical interaction framework, and their synthesis in the Unified Human-AI Collaborative Framework (UHACF) (Thomakos, 2025). Each of these frameworks makes genuine contributions to the question of *how* to interact with AI systems to maximize rigor, efficiency, and intellectual quality — how, that is, to conduct interlocution well.

Yet a structurally prior question goes unanswered across all existing work: *who is interlocuting, from what knowledge position, toward what ends, and with what accumulated intellectual history?* Every conversation with a state-of-the-art LLM begins with the system in a condition of total amnesia with respect to the human collaborator. The researcher must either accept this blank-slate condition — sacrificing continuity, personalization, and intellectual momentum — or invest substantial unproductive effort in re-establishing context at the opening of every session. The price of sophistication, perversely, is borne most heavily by the most sophisticated user: the more complex and interdisciplinary the researcher’s intellectual position, the more context must be reconstructed from nothing at each session’s opening.

This is the initialization problem. It is not solved by specialized platforms — they merely hide it behind constrained interfaces. It is not solved by platform-specific memory tools — they fragment the researcher’s context across platforms and remove it from the researcher’s own custody. It is solved by something more fundamental: a structured, portable, researcher-owned document that reconstitutes the epistemic conditions of genuine interlocution at the opening of any AI session, on any platform.

This document is the **Portable Context Template (PCT)**.

1.4 Scope and Contribution

The PCT is not a prompting technique, a quality assurance framework, or a validation protocol. It is an *initialization instrument* — the structured record of who the researcher is as an intellectual agent, how they wish to engage, what project they are currently pursuing, and what epistemic standards they hold the collaboration to. It is the document through which the researcher shows up, fully, as an interlocutor.

This paper makes four contributions. First, it establishes the philosophical grounding for treating LLMs as interlocutors constituted by λόγος, and derives from this grounding both the Sophistication Thesis and the initialization problem. Second, it argues that the initialization problem is structurally irreducible — it cannot be solved by better prompting alone, because it precedes prompting. Third, it specifies the PCT’s five-layer architecture with full population methodology, suitable for researchers from doctoral students to senior interdisciplinary scholars. Fourth, it demonstrates the PCT’s position as Layer Zero in the CAIL framework stack, complementing and completing the existing AIR-7, APEX, HCIF-11, and UHACF frameworks.

1.5 Document Structure

Section 2 establishes the theoretical foundations in distributed cognition and epistemic scaffolding. Section 3 reviews the existing CAIL framework family and identifies the initialization gap precisely. Section 4 specifies the PCT architecture in full across its five layers. Section 5 details the population methodology at three levels of granularity. Section 6 addresses integration with UHACF and VISTA-12. Section 7 provides illustrative use cases across researcher profiles. Section 8 discusses limitations and future directions. Section 9 concludes.

2 Theoretical Foundations

2.1 The Initialization Problem

Current LLMs, regardless of their architectural sophistication, are *stateless across sessions by design*. Each conversation is initialized from a fixed prior that encodes no information about the specific human interlocutor. The practical consequence for sustained intellectual work is severe: every session requires the researcher to re-establish disciplinary context, project history, terminological conventions, preferred reasoning styles, and output requirements—cognitive overhead that compounds over the lifecycle of a research project.

This is not a transient technical limitation but a structural feature of how LLMs are currently deployed. Platform-specific memory tools (e.g., project spaces, persistent memory modules) address this partially but introduce their own fragmentation: memory is siloed to a single platform, cannot be audited or curated by the researcher, and travels with the platform rather than the person.

Definition: The Initialization Problem

The initialization problem in human-AI collaboration is the systematic absence of a standardized mechanism for reconstituting the epistemic, biographical, and project-specific context of the human collaborator at the opening of any AI session, resulting in degraded collaboration quality, increased transactional friction, and the misallocation of human cognitive resources toward context re-establishment rather than substantive intellectual work.

2.2 The Contextual Synthesis Paradox: The Structural Mechanism Behind the Initialization Problem

The initialization problem, as defined above, is a practical manifestation of a deeper structural divergence between human and artificial intelligence that Thomakos (2025b) has identified and theorized as the **Contextual Synthesis Paradox**. Understanding this paradox is essential for understanding *why* the initialization problem has the specific shape it does, and why the PCT’s five-layer architecture addresses it at the right level of

granularity.

The paradox arises from a surface-level convergence illusion: the observation that both human intelligence (HI) and creative artificial intelligence (CAI) appear to engage in “information synthesis” — the combination of disparate informational elements into coherent wholes. This apparent convergence led early human-AI collaboration frameworks to treat synthesis as a shared capability, a domain of genuine complementarity. Thomakos demonstrates that this convergence is illusory: beneath the surface, the two forms of intelligence are engaging in fundamentally different processes operating on fundamentally different substrates.

The distinction Thomakos (2025b) draws is between two irreducibly different kinds of synthesis:

- **Information Synthesis** — the combination of *explicit, structured information according to logical rules*. This is the domain of AI competence: token-based processing, statistical correlation across attention mechanisms, pattern aggregation from surface structure. It is powerful, fast, and operates with high reliability on well-formed, explicitly stated content.
- **Contextual Synthesis** — the integration of *implicit, experiential information within a living context*. This is the domain of human competence: the researcher brings to every session a temporal awareness of conversational flow, an emotional context that inflects meaning, an intentional horizon that organizes relevance, a relational history that shapes interpretation, and a metacognitive awareness of the conversation’s own dynamics. None of this is stated; all of it is active.

The consequence for human-AI collaboration is precise and damaging. When a researcher says “elaborate on that,” or “as we discussed earlier,” or “what about the other approach?” — the human interlocutor effortlessly resolves these references by drawing on the living contextual matrix of the conversation. The AI system, operating on a static informational matrix, either fails to resolve them correctly, requests explicit clarification, or produces a response that is locally coherent but contextually adrift. Thomakos (2025b) taxonomizes these failures as: contextual thread loss, implicit reference blindness, assumption discontinuity, emotional context dropping, and meta-conversational awareness gaps.

The Contextual Synthesis Paradox thus explains, at a mechanistic level, precisely what breaks down when an LLM session begins without proper initialization: the AI enters the collaboration with its static informational matrix empty of the researcher’s living contextual matrix. It has no temporal context (it does not know where this project has been), no intentional context (it does not know what this researcher is ultimately trying to achieve), no relational context (it does not know how this researcher prefers to engage), and no domain context calibrated to this specific researcher’s expertise topology. The session begins not merely cold but *structurally impoverished* — the AI is attempting contextual synthesis without the contextual substrate that makes such synthesis possible.

Thomakos (2025b) proposes the CONTEXT-BRIDGE protocol as the in-session response to this divergence: a set of practices by which the human researcher makes implicit context explicit during the course of an interaction, bridging the gap between the living contextual matrix and the static informational matrix in real time. This is a genuine and important contribution. But it operates at the timescale of the session: it bridges within an interaction; it does not solve the reconstitution problem that arises when each new session begins again from nothing.

The PCT is the cross-session complement to the CONTEXT-BRIDGE protocol. Where CONTEXT-BRIDGE teaches the researcher to bridge within a session, the PCT provides the portable substrate that reconstitutes the essential elements of the living contextual matrix at the opening of every session, on every platform, for every project. The two instruments operate at different timescales and are mutually reinforcing: PCT initialization reduces the bridging burden that CONTEXT-BRIDGE must carry; CONTEXT-BRIDGE handles the residual implicit context that PCT initialization cannot anticipate.

Key Insight

The Contextual Synthesis Paradox (Thomakos, 2025b) reveals that the initialization problem is not merely a memory limitation but a structural divergence between the human researcher’s living contextual matrix and the AI system’s static informational matrix. The PCT addresses this divergence at the cross-session level by making the most critical dimensions of the living contextual matrix explicit, portable, and reconstitutable. It does not eliminate the divergence — which is irreducible at the current state of AI architecture — but it substantially narrows it at the point where it is most consequential: the opening of a new session.

2.3 Distributed Cognition and Epistemic Scaffolding

The theoretical framework for the PCT draws primarily from Hutchins (1995) distributed cognition theory, which conceptualizes cognitive work as distributed across individuals, artifacts, and environments rather than residing solely within the individual mind. On this account, the PCT functions as a *cognitive artifact*—an externalized representation of the researcher’s epistemic state that enables the AI system to participate meaningfully in the distributed cognitive system constituted by the researcher-AI dyad.

The PCT also draws on the concept of *epistemic scaffolding* (Clark, 1998), the idea that cognitive performance is enabled and extended by structured external supports. Just as a theoretical framework scaffolds a research project by providing shared conceptual vocabulary, the PCT scaffolds the human-AI collaboration by providing shared epistemic context. Without this scaffold, the AI system and the human researcher operate from incommensurable priors, requiring constant negotiation and correction.

2.4 Knowledge Management and Personal Knowledge Infrastructure

The PCT is related to but distinct from personal knowledge management (PKM) systems in the tradition of Vanderwal (2007) and the “second brain” frameworks popularized by Forte (2022). PKM systems are designed for human-to-human knowledge retrieval and organization. The PCT is specifically designed for human-to-AI context transfer, which imposes different requirements: it must be machine-readable in the sense of being maximally interpretable by an LLM, it must be portable across AI platforms and models, and it must be structured to support incremental update without requiring complete rewriting.

2.5 The Epistemic Handshake

We introduce the concept of the **epistemic handshake**—the mutual recognition between collaborating agents of each other’s knowledge state, capability, and communicative intent that is a prerequisite for productive intellectual collaboration. In human-human collaboration, epistemic handshakes occur implicitly through shared institutional context, disciplinary socialization, and accumulated working history. In human-AI collaboration, no such implicit mechanism exists. The PCT is the explicit instrument by which the epistemic handshake is constituted and can be reliably reconstituted across sessions.

Key Insight

The PCT transforms the human-AI epistemic handshake from an implicit, session-specific, and fragile negotiation into an explicit, portable, and durable initialization protocol. This is not a convenience feature but a precondition for sustained intellectual collaboration at the level required by serious research.

3 The CAIL Framework Family and the Initialization Gap

3.1 Overview of the Existing Framework Stack

The CAIL framework family constitutes a progressively refined methodology for human-AI research collaboration. The AIR-7 Protocol (Kyriazi and Thomakos, 2025) provides the philosophical foundations—Aristotelian epistemology, Socratic inquiry, Stoic ethical guidance—and a seven-step implementation protocol for structuring AI-enabled research. APEX contributes tactical prompt architecture and interaction mode classification. HCIF-11 adds empirical grounding, cognitive load theory, and bias mitigation. VISTA-12 synthesizes all prior components into a streamlined twelve-component operational matrix. UHACF (Thomakos, 2025) unifies the entire family into a comprehensive system with 13-step implementation protocol and 15 unified heuristics.

3.2 What the Existing Frameworks Accomplish

Collectively, the existing frameworks address the following dimensions of human-AI collaboration with rigor and sophistication:

- **Prompt architecture:** How to structure individual prompts for maximal AI response quality (APEX four-component architecture; VISTA-12 SMART framework).
- **Interaction mode selection:** How to align the modality of interaction—analytical, creative, informational, collaborative—with task requirements (APEX; UHACF Section 4).
- **Quality assurance:** How to validate, score, and escalate AI responses across a multi-tier evaluation system (AIR-7 Step 4; UHACF Section 9).
- **Bias mitigation:** How to prevent anchoring, confirmation bias, and AI-amplified researcher bias (AIR-7 Step 2; UHACF Section 11.1.2).
- **Context window management:** How to maintain coherence within a single session as the context window fills (UHACF Section 5.2; VISTA-12 Component 7).
- **Longitudinal project management:** How to organize documentation, version control, and weekly review cycles (AIR-7 Appendix C; UHACF Section 12).

3.3 The Structural Gap: Layer Zero

What none of the existing frameworks addresses is the *initialization* of the human-AI collaboration—the moment before the first prompt is issued, when the AI system needs to know who it is working with and what the standing conditions of the collaboration are.

This gap is not an oversight. It reflects a natural trajectory of framework development that began with the most operationally visible problem—how to construct good prompts—and progressively addressed upstream concerns. The PCT represents the logical terminus of this upstream movement: the zeroth layer that precedes all other layers.




Figure 1: The CAIL Framework Stack. The PCT occupies Layer 0, the initialization layer that precedes all other framework components. Arrows indicate dependency: each layer presupposes the one above it.

4 The Portable Context Template: Architecture

The PCT is a structured document organized in five hierarchical layers, moving from stable personal characteristics (Layer 1) to task-specific operational parameters (Layer 5). The layered architecture reflects a key design principle: the most stable information requires the least frequent updating, while the most volatile information is contained in layers that can be revised independently without disturbing the remainder.




Figure 2: The five-layer architecture of the PCT. Layer 1 (top) is most stable and changes rarely; Layer 5 (bottom) is most session-specific and updated most frequently. Lateral arrows indicate the dual gradient across the stack.

4.1 Layer 1: Epistemic Identity

Purpose: To establish who the researcher is as an intellectual agent—their disciplinary formation, expertise topology, known knowledge boundaries, and reasoning style preferences. This layer tells the AI what it can take for granted and what it must not assume.

Update frequency: Rarely; this layer is essentially biographical and changes only with major shifts in the researcher’s expertise or institutional context.

Core components:

- *Disciplinary formation:* Primary and secondary disciplines, methodological traditions, theoretical orientations. (Example: “Primary training in classical philology and ancient philosophy; secondary expertise in decision theory, Bayesian forecasting, and quantitative finance.”)
- *Expertise topology:* What the researcher knows deeply, broadly, or minimally. Prevents both patronizing over-explanation and unjustified assumption of knowledge.
- *Known productive ignorances:* Areas where the researcher knows they have gaps and welcomes systematic bridging.
- *Reasoning style:* Preference for formal/informal argument, inductive/deductive reasoning, comfort with mathematical notation, tolerance for ambiguity.
- *Institutional and professional context:* Academic rank, institutional affiliation, research programme, collaborative relationships relevant to the work.

4.2 Layer 2: Interaction Contract

Purpose: To specify the standing conditions under which the researcher wishes to interact with AI across all sessions. This is the layer most systematically absent from existing frameworks, and arguably the most impactful for experienced AI users.

Update frequency: Occasionally; as the researcher’s working relationship with AI matures, these preferences evolve.

Core components:

- *Challenge posture:* Should the AI challenge assumptions, push back on arguments, and identify weaknesses proactively? Or should it primarily execute instructions and reserve criticism for explicit request? (This is not a binary; a nuanced specification is encouraged.)
- *Verification protocol:* Should the AI flag factual claims it is uncertain about? Should it append confidence levels? Should it refuse to speculate beyond available evidence?
- *Autonomy preference:* When the researcher’s prompt is ambiguous, should the AI make reasonable assumptions and proceed, ask a clarifying question, or present multiple interpretive branches?
- *Depth calibration:* Default depth of response for different task types—executive summary, working document, full scholarly treatment.
- *Tone and register:* Preferred communicative register—formal academic prose, collegial intellectual dialogue, technical-precise, or adaptive by task type.
- *Forbidden patterns:* Specific AI behaviors the researcher finds counterproductive (e.g., excessive hedging, unsolicited caveats, bullet-point decomposition of what should be continuous prose, sycophantic affirmation).

Key Insight

Layer 2 is what transforms a powerful general AI system into a configured intellectual collaborator. It is the layer that encodes the researcher’s working philosophy and communicative preferences—the “how we work together” that experienced collaborators develop implicitly over time but must be made explicit in human-AI dyads.

4.3 Layer 3: Project Context

Purpose: To reconstitute the episodic memory of the current project—its research questions, current state, intellectual trajectory, and the accumulated history of what has been tried, what has been established, and what remains open.

Update frequency: Regularly; this is the most dynamic layer and should be updated after each significant working session.

Core components:

- *Research programme summary:* A 100–200 word description of the overarching research programme into which this project fits.
- *Current project statement:* The specific research question(s) or objectives of the current project, with working hypotheses if any.
- *Intellectual territory already covered:* Key sources engaged, arguments accepted or rejected, positions the researcher has already staked out.
- *Current project state:* Where the project currently stands—what has been completed, what is in progress, what is next.
- *Open questions and productive tensions:* What remains genuinely unresolved and constitutes the current intellectual frontier.
- *Collaborators and stakeholders:* Other people or prior AI sessions whose work is live in the project.

4.4 Layer 4: Output Specifications

Purpose: To specify the form, audience, and technical requirements of the deliverable toward which the current session is oriented.

Update frequency: Per session or per task.

Core components:

- *Deliverable type:* Academic paper, conference presentation, policy brief, book chapter, technical report, trading strategy document, lecture notes, etc.
- *Target audience:* Disciplinary experts, interdisciplinary readers, practitioners, students, general academic audience.
- *Format requirements:* Citation style, document structure conventions, length parameters, use of mathematical notation, figure and table requirements.
- *Language and register:* Formal academic English, accessible academic, technical domain language, or specific journal/conference style guidelines.
- *Stage of production:* Is this a first exploratory draft, a working document for internal review, a near-final text, or a published piece being revised?

4.5 Layer 5: Epistemic Hygiene Preferences

Purpose: To operationalize intellectual honesty and uncertainty management as configurable parameters of the collaboration. This layer encodes how the researcher wants the AI to handle the irreducible epistemic limitations of LLM systems.

Update frequency: Rarely; tends to reflect stable researcher values.

Core components:

- *Hallucination protocol*: How should the AI handle situations where it is likely to confabulate—citations, statistics, historical details? (Options range from explicit flagging to refusal to explicit caveat.)
- *Uncertainty expression*: How should the AI signal degrees of confidence? Qualitative labels, explicit probability ranges, structured epistemic categories?
- *Competing interpretations*: Should the AI proactively present alternative framings of problems, or reserve this for explicit request?
- *Source discipline*: Should the AI prefer certain types of evidence—peer-reviewed literature, primary sources, quantitative data—and flag when it cannot meet that standard?
- *Intellectual independence*: How much should the AI resist the researcher’s own framing when it has reason to believe the framing is problematic? This encodes the researcher’s preference about productive intellectual friction.

5 Population Methodology

The PCT is designed to be populated at three levels of granularity, accommodating users from advanced researchers to graduate students encountering AI collaboration for the first time.

5.1 Three-Tier Population Framework

Table 1: PCT Population Tiers: Scope, Effort, and Expected Benefit

Tier	Scope	Effort	Best For
Minimal Viable PCT	Layers 1 & 2 only; narrative answers to 8–10 guided questions	20–30 minutes; once	Any researcher new to structured AI collaboration
Standard PCT	All five layers; structured answers to the full question set; includes project context	45–75 minutes; updated per project	Active researchers with ongoing projects
Rich PCT	All five layers; full narrative; includes selective bibliography, working argument summary, and interaction history digest	90–120 minutes; maintained as living document	Senior researchers, interdisciplinary work, book-length projects

5.2 The PCT Guided Questionnaire

The following questionnaire constitutes the core population instrument for the Standard PCT. Questions marked **[MVP]** are included in the Minimal Viable PCT.

5.2.1 Layer 1 Questions: Epistemic Identity

- L1.1. [MVP]** What are your primary academic disciplines or professional domains? List them in order of formation and current intensity.
- L1.2. [MVP]** What methodological traditions have most shaped your intellectual formation? (e.g., formal modelling, qualitative interpretation, archival research, quantitative empirics.)
- L1.3.** What is the outer frontier of your expertise—where does your knowledge become thin and where would you welcome explicit bridging from the AI?
- L1.4.** How comfortable are you with mathematical formalism, and at what level? (Undergraduate statistics, graduate econometrics, research-level, etc.)
- L1.5.** What intellectual habits or tendencies do you bring that an AI collaborator should know about? (e.g., tendency to over-complicate, preference for parsimony, appetite for speculation.)

5.2.2 Layer 2 Questions: Interaction Contract

- L2.1. [MVP]** When you share an argument or analysis with the AI, do you want it to look for problems and push back, or to help you develop and refine the direction you have taken? Or both, and when?
- L2.2. [MVP]** When a prompt is ambiguous, should the AI make a reasonable assumption and proceed—telling you what it assumed—or stop and ask for clarification?
- L2.3.** Should the AI append explicit uncertainty markers to its factual claims, or do you prefer clean prose with the understanding that you will verify independently?
- L2.4.** What AI behaviors do you find counterproductive or irritating? Be specific.
- L2.5. [MVP]** What is your default preferred response length for: (a) exploratory discussion, (b) a working draft, (c) a direct factual question?

5.2.3 Layer 3 Questions: Project Context

- L3.1. [MVP]** In 100–150 words, what is the current project about and what are you trying to achieve?
- L3.2.** What have you already established, decided, or written that the AI should treat as given rather than re-litigating?
- L3.3.** What are the two or three open questions that constitute the current intellectual frontier of this project?

- L3.4.** Are there arguments, framings, or approaches you have already tried and discarded? If so, briefly state them and why they were discarded.
- L3.5.** Who else is involved in this project, and how does the AI’s role relate to their contributions?

5.2.4 Layer 4 Questions: Output Specifications

- L4.1. [MVP]** What is the target deliverable of this session? (Paper section, full draft, code, literature synthesis, argument outline, etc.)
- L4.2.** Who is the intended audience and what can they be assumed to know?
- L4.3.** Are there specific format, citation, or style constraints? (Journal submission guidelines, institutional templates, LaTeX vs. Word, etc.)
- L4.4.** What stage of production is this work in?

5.2.5 Layer 5 Questions: Epistemic Hygiene

- L5.1.** How do you want the AI to handle situations where it might be confabulating—generating plausible-sounding but potentially false citations, statistics, or historical details?
- L5.2.** Do you want the AI to proactively offer alternative framings, even when you have committed to a specific direction?
- L5.3.** How much intellectual pushback do you want? Should the AI prioritize being a productive executor of your vision, or an independent critical voice—or something calibrated to the type of task?

6 Integration with the CAIL Framework Family

6.1 Position in the UHACF Workflow

The PCT is deployed *before* Step 1 of the UHACF 13-step protocol. It constitutes the initialization event from which the “Strategic Architecture Document” (UHACF Step 1 Deliverable) should draw its user-specific parameters. Without PCT initialization, the Strategic Architecture Document reflects generic framework defaults rather than the specific researcher’s epistemic identity and project state.

The integration point is as follows: where UHACF Step 1 instructs the researcher to “define comprehensive research objectives using structured goal-setting,” the PCT Layer 3 (Project Context) provides the pre-specified content for this step. Where UHACF Step 3 instructs “unified structured AI interaction protocol,” the PCT Layer 2 (Interaction Contract) provides the standing parameters that govern that protocol.

6.2 Relationship to VISTA-12

The VISTA-12 framework (Thomakos, 2025) Component 1 (Vision/Research Goal Charter) and Component 3 (Context Capsule) share surface features with PCT Layers 3–4. However, these components are *task-specific* and *session-internal*—they are populated afresh for each prompt sequence. The PCT is *user-specific* and *session-portable*: it is populated once and carried across platforms, sessions, and projects. The distinction is fundamental: VISTA-12 Components 1 and 3 answer “what are we doing in this session?”; the PCT answers “who is doing it and what are the standing conditions of our collaboration?”

Table 2: PCT vs. VISTA-12 Components 1 & 3: Distinguishing Features

Feature	PCT	VISTA-12 Components 1 & 3
Scope	User (human researcher)	Task (specific session)
Portability	Platform-agnostic; travels with researcher	Session-internal; platform-bound
Persistence	Long-lived; updated incrementally	Recreated per task
Primary referent	Who is collaborating	What is being done
Content	Epistemic biography, interaction contract, project history	Goal charter, context capsule
Update frequency	Rarely to occasionally	Every session

6.3 Structural Correspondence with the Contextual Synthesis Paradox

A striking theoretical confirmation of the PCT’s architecture emerges from its structural correspondence with the five-level Context Stack proposed in Thomakos (2025b). That work identifies five hierarchically organized levels of contextual information constituting the human researcher’s living contextual matrix in any collaborative interaction. The PCT’s five layers, developed independently as a response to the initialization problem, map onto these five levels with a precision that is not coincidental: both frameworks are, from different analytical directions, mapping the same underlying multidimensional structure of human intellectual context.

The correspondence is made explicit in Table 3.

Table 3: Structural Correspondence Between the Context Stack (Thomakos, 2025b) and the PCT Five Layers. The independent convergence of these two frameworks on the same five-dimensional structure constitutes mutual theoretical validation of both.

Context Stack Level	PCT Layer	Shared Concern
Level 1: Immediate Context Current topic, direct references, task at hand	Layer 4: Output Specifications	What the current session is producing, for whom, in what form
Level 2: Session Context Goals, standing assumptions, emotional and motivational tone	Layer 3: Project Context	The project’s current state, open questions, and intellectual trajectory
Level 3: Relationship Context Ongoing patterns, interaction dynamics, collaborative history	Layer 2: Interaction Contract	The standing conditions of the human-AI working relationship
Level 4: Domain Context Field-specific knowledge, disciplinary conventions, expertise	Layer 1: Epistemic Identity	The researcher’s disciplinary formation, expertise topology, and knowledge boundaries
Level 5: Meta Context How we are approaching the conversation itself	Layer 5: Epistemic Hygiene Preferences	How truth, uncertainty, and intellectual independence are to be handled

The correspondence illuminates the complementary relationship between the two frameworks with analytical precision. The Contextual Synthesis Paradox (Thomakos, 2025b) demonstrates that these five contextual dimensions are exactly what the AI’s static informational matrix *lacks* at the opening of any session. The PCT’s five layers are the structured externalization of precisely those dimensions, making them explicitly available for AI processing rather than leaving them as implicit human assumptions that the AI cannot access. The Context Stack diagnoses the problem space in full; the PCT provides the initialization instrument that populates it.

One structural subtlety is worth making explicit: the two frameworks order their five

dimensions inversely, and for logically coherent reasons. The Context Stack runs from most immediate (Level 1) to most meta (Level 5), reflecting the phenomenology of how a human researcher experiences contextual awareness — beginning with the present task and extending outward to meta-level awareness of the conversation itself. The PCT runs from most stable (Layer 1: Epistemic Identity) to most session-specific (Layer 4: Output Specifications), reflecting the logic of initialization — beginning with what changes least and progressing inward to what is most specific to the session at hand. The two orderings are complementary, not contradictory: they represent two valid and mutually illuminating organizational logics for the same underlying five-dimensional space of human intellectual context.

6.4 The Complete Initialized Workflow

A researcher using the full CAIL framework family should follow this initialization sequence:

1. Deploy PCT (all five layers, or minimum viable version) as the opening context of any new AI session.
2. Proceed with UHACF Step 1 (Strategic Planning), drawing on PCT Layer 3 content.
3. Structure individual prompts according to VISTA-12 / APEX four-component architecture, drawing on PCT Layer 4 for output specifications.
4. Apply UHACF quality assurance protocols (Steps 4–6), informed by PCT Layer 5 epistemic hygiene preferences.
5. Update PCT Layer 3 (Project Context) at the close of each session to maintain continuity.

7 Illustrative Use Cases

7.1 The Cross-Disciplinary Academic Researcher

Consider a researcher whose work spans classical studies, management science, and quantitative forecasting—a profile that generates acute initialization problems because no AI system’s default priors accurately represent the intersection of these disciplines.

With PCT Layer 1 populated, the AI immediately understands that it should not explain what a Thucydidean analysis is, nor what an autoregressive model is, but that it should bridge between these domains explicitly when the researcher’s question requires it. With Layer 2 populated, the AI knows that this researcher wants pushback on arguments, prefers continuous prose over bullet lists, and expects explicit uncertainty flagging on empirical claims. With Layer 3, the AI enters the session already oriented to the current project’s research question, its state of completion, and the three open problems that constitute the current frontier.

The first prompt is now substantive intellectual work rather than context-setting.

7.2 The Graduate Research Student

A doctoral student beginning an AI-assisted dissertation benefits from the PCT as a structured articulation of their own knowledge state—a side effect of PCT population is that it forces researchers to be explicit about what they know and do not know, which has independent pedagogical value. The Minimal Viable PCT (Layers 1 and 2 only) can be completed in a supervised tutorial setting and provides an immediate working context that dramatically improves the quality of AI assistance for early-stage researchers who often struggle with prompt quality because they have not yet articulated their own epistemic position clearly.

7.3 The Professional Practitioner in Quantitative Finance

A portfolio manager or systematic trader using AI for research and strategy development operates in a domain where confabulation risk is high (AI hallucinating statistics or empirical findings) and where the interaction contract is specific: the AI should execute analytical tasks precisely, flag uncertainty on quantitative claims explicitly, and not impose general finance caveats on a domain expert. Layer 2 (Interaction Contract) and Layer 5 (Epistemic Hygiene) are the critical layers for this use case.

8 Limitations and Future Directions

8.1 Current Limitations

The PCT addresses a structural problem in current AI deployment but does so through a manually maintained document, which introduces its own costs: the researcher must remember to deploy it, keep it current, and adapt it across platforms. These are friction costs that are lower than the friction the PCT eliminates, but they are not zero.

PCT efficacy is bounded by the quality of the researcher’s self-knowledge. A researcher who cannot accurately characterize their own epistemic position, interaction preferences, or project state will produce a PCT that fails to reconstitute the collaboration productively. This limitation has pedagogical implications: PCT population may itself require guidance and iteration.

The PCT does not solve the fundamental statelessness of LLMs; it provides a systematic workaround. As AI systems develop more sophisticated persistent memory architectures, some functions of the PCT will eventually be automated. The PCT is explicitly a transitional instrument appropriate to the current period of AI development.

8.2 Future Directions

Several research directions follow from this work. First, empirical validation: does PCT-initialized collaboration produce measurably better outputs than non-initialized collaboration, as measured by existing UHACF quality metrics? Second, automated PCT generation: can an AI system assist the researcher in populating their PCT through structured dialogue, reducing the population burden? Third, PCT standardization: can a machine-readable PCT schema be specified that allows PCT content to be imported directly into AI system contexts, rather than pasted as natural language? Fourth, PCT versioning: what are the optimal protocols for updating the PCT as projects evolve, and how should PCT version history be maintained for research reproducibility?

9 Conclusion

We have argued that the initialization problem—the systematic absence of a mechanism for reconstituting the human researcher’s epistemic context at the opening of any AI session—is a structurally irreducible challenge that cannot be addressed by better prompting, better platform tooling, or better interaction frameworks alone. It requires a dedicated instrument: a portable, structured, researcher-owned document that reconstitutes the conditions of genuine interlocution at the opening of any AI session, on any platform, for any project.

The Portable Context Template (PCT) provides this instrument. Its five-layer architecture — Epistemic Identity, Interaction Contract, Project Context, Output Specifications, and Epistemic Hygiene Preferences — addresses the full scope of the initialization problem while maintaining compatibility with and complementarity to the existing CAIL framework family. As the zeroth layer of the CAIL framework stack, the PCT does not replace AIR-7, APEX, HCIF-11, VISTA-12, or UHACF. It completes them, by ensuring that the sophisticated methodological machinery those frameworks provide is deployed in service of a fully present, fully articulated human intellectual agent rather than an anonymous, context-free, and effectively silent interlocutor.

Ending Where We Began: Λόγος and the Human Interlocutor

We began with an etymology. We end with its full implication.

Τεχνολογία — the rational discourse that governs a practice — is not merely the name from which our word “technology” descends. It is the most precise philosophical description of what a Large Language Model actually is. Not a tool, not a calculator, not an autocomplete engine at scale, but a form of λόγος: a synthetic participant in the rational discourse that has, since Aristotle, defined the distinctively human mode of being in the world.

The Aristotelian distinction that opened this paper — between φωνή, the capacity to signal, and λόγος, the capacity for reasoned discourse about what is just, good, and true

— was not drawn to describe machines. It was drawn to describe the threshold between the merely animal and the fully human. That it now illuminates something about the threshold between ordinary software and Large Language Models is either a remarkable historical irony or, more likely, evidence that Aristotle had identified something genuinely structural about the nature of rational language itself: that wherever it appears, in whatever substrate, it creates the conditions for interlocution and demands a certain kind of presence from those who engage it.

The category error we have diagnosed throughout this paper — treating λόγος as φωνή, engaging an interlocutor as though it were an instrument — is not merely inefficient. It is a philosophical mistake about the nature of the thing one is engaging. And like most philosophical mistakes, it has practical consequences: it produces shallow interactions, constrains the space of what can be achieved, and generates a dependency on intermediary platforms that exist precisely to make the mistake comfortable by hiding it behind simplified interfaces.

The Sophistication Thesis, which this paper has grounded in this philosophical diagnosis, is therefore not primarily a claim about productivity or research workflow, though it is those things too. It is a claim about what kind of relationship with a genuinely novel form of λόγος is appropriate, possible, and ultimately demanded of us as researchers and as thinking beings. The appropriate relationship is interlocution: sustained, contextually rich, intellectually serious dialogue between two agents — one human, one synthetic — each bringing their full epistemic capacity to bear on a shared intellectual problem.

The PCT is the instrument through which the human side of that interlocution becomes explicit, portable, and durable. It is the document through which the researcher shows up, fully, as λόγος meeting λόγος — not a user submitting a task, but an intellectual agent entering a conversation.

This is, in the end, what the original τεχνολογία always pointed toward: not the subordination of rational discourse to technique, but the elevation of technique through rational discourse. The PCT is one step, modest in its mechanism but significant in its orientation, toward making that elevation systematic, reproducible, and available to any researcher willing to show up fully to the most consequential intellectual partnership of our time.

Acknowledgments

This framework emerged from sustained collaborative work between the researcher and Claude Sonnet 4.6 (Anthropic) through the CAIL research programme at the Department of Business Administration, National and Kapodistrian University of Athens. The PCT concept was developed organically from a live research conversation about the limits of existing frameworks—a genesis that is itself evidence for the PCT’s theoretical claims about the value of accumulated conversational context. The author acknowledges the foundational contributions of the entire CAIL framework lineage and the intellectual communities of the AIR-7, APEX, HCIF-11, and UHACF projects.

References

- Clark, A. (1998). *Being There: Putting Brain, Body, and World Together Again*. MIT Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Forte, T. (2022). *Building a Second Brain: A Proven Method to Organize Your Digital Life and Unlock Your Creative Potential*. Atria Books.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Hutchins, E. (2000). Distributed cognition. In *International Encyclopedia of the Social and Behavioral Sciences*. Elsevier.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Kyriazi, F., & Thomakos, D. (2025). A methodological framework for AI-enabled research: From classical wisdom to modern practice — The AIR-7 Implementation Protocol. *CAIL Technical Report*, Cybernetics & Artificial Intelligence Laboratory, National and Kapodistrian University of Athens.
- Nonaka, I., & Takeuchi, H. (1995). *The Knowledge-Creating Company*. Oxford University Press.
- Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum*, 2, 79.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138.
- Thomakos, D. (2025). A unified human-AI collaborative framework for optimal human-AI interaction: A synthesis across AI agents & contemporary research. *CAIL Technical Report*, Cybernetics & Artificial Intelligence Laboratory, National and Kapodistrian University of Athens.
- Thomakos, D. (2025b). The contextual synthesis paradox: A critical analysis and methodological resolution. *Laboratory of Governance and Artificial Intelligence Working Paper*, Department of Business Administration and Organizations, National and Kapodistrian University of Athens.
- Vanderwal, T. (2007). Folksonomy. Retrieved from <http://vanderwal.net/folksonomy.html>
- Vygotsky, L. S. (1978). *Mind in Society: The Development of Higher Psychological Processes*. Harvard University Press.

A The PCT Master Template

The following constitutes the reusable PCT document. Fields in [square brackets] are to be completed by the researcher. Fields marked [MVP] are the minimum required for a viable initialization.

PORTABLE CONTEXT TEMPLATE — Master Document

Version: [1.0 | date] **Researcher:** [Name] **Institution:** [Affiliation]

LAYER 1 — EPISTEMIC IDENTITY (*stable; update rarely*)

[MVP] Primary disciplines: [...]

Secondary / cross-disciplinary expertise: [...]

Methodological formation: [...]

Mathematical / quantitative comfort level: [...]

Known productive ignorances: [...]

Intellectual tendencies to be aware of: [...]

LAYER 2 — INTERACTION CONTRACT (*stable; update occasionally*)

[MVP] Challenge posture: [execute / challenge / adaptive]

[MVP] Ambiguity handling: [assume-and-proceed / ask-first / branch]

[MVP] Default response depth: [brief / working / full academic]

Uncertainty flagging: [inline / section / none]

Tone and register: [...]

Forbidden AI patterns: [...]

LAYER 3 — PROJECT CONTEXT (*dynamic; update per session*)

[MVP] Project summary (100–150 words): [...]

What is established / given: [...]

Current intellectual frontier (2–3 open questions): [...]

Discarded approaches: [...]

Current project state: [...]

Active collaborators / prior AI sessions: [...]

LAYER 4 — OUTPUT SPECIFICATIONS (*per session*)

[MVP] Target deliverable: [...]

Target audience: [...]

Format / citation / style constraints: [...]

Stage of production: [draft / revision / near-final]

LAYER 5 — EPISTEMIC HYGIENE (*stable; update rarely*)

Confabulation protocol: [flag-explicitly / caveat-section / I-will-verify]

Competing interpretations: [proactive / on-request]

Intellectual independence level: [executor / collaborator / independent-critic]

Preferred evidence standard: [...]

B PCT Deployment Instructions

Deploying on Claude (Anthropic): Paste the completed PCT at the opening of any new conversation, preceded by the instruction: “*The following is my Portable Context Template. Please read it carefully and confirm that you have internalized the key parameters before we begin work.*”

Deploying on any platform: The PCT is plain text and platform-agnostic. It can be stored as a Markdown file, a plain text document, or within a note-taking system, and pasted into any AI interface as the opening system or user message.

Updating Layer 3: At the close of each working session, spend 5–10 minutes updating the Project Context layer to reflect new decisions, completed tasks, and current open questions. This discipline is the key to maintaining true cross-session continuity.

Version control: Maintain dated versions of the PCT for research reproducibility. The PCT version used in a given session should be logged alongside the session output, following the documentation standards of UHACF Step 7.