IADIS International Journal on Computer Science and Information Systems
Vol. 21, No. 2, pp. 18-32
ISSN: 1646-3692

FROM THEORIES TO TOOLS: THE *AD HOC WEAVING FRAMEWORK* WITH AN AI AMANUENSIS FOR INNOVATION PRACTICE

Mario Bonatti[1], Mike Martin[2] and Gianni Jacucci[3]
[1]*Retired from Italtel spa, Milan, Italy*
[2]*Northumbria University, Newcastle, UK*
[3]*University of Trento, Trento, Italy*

ABSTRACT

The Ad Hoc Weaving Framework (AHWF) was introduced as a design-science meta-artifact that turns deep commitments about ontology, value, agency and communication into practically usable tools for innovation management, entrepreneurship and technology governance (Bonatti et al., 2025). It provides five tightly interwoven pillars – a trope-based meta-ontology with explicit life-cycles and time-modes; a layered distinction between natural processes and human processings; a unified account of persons and organisations as Selfs' Complexes; a Value–Life Complex (VLC) that ties mission to governance, learning, planning, doing and facts; and a multi-level communications pragmatics (CI–CV) supported by an engineering stack that shifts attention from message transmission to completion – together with a Minimal Working Set (MWS) of three one-page instruments that teams can run in practice. At the same time, large language models have matured to the point where they can serve as AI amanuenses: persistent conversational companions that can read and write across many documents, maintain a memory of a project's evolving conceptual weave, detect inconsistencies, recover forgotten commitments, and prompt reflexive re-organisation. This paper explores how AHWF and an AI amanuensis can be combined into a second-order meta-artifact for innovation: AHWF provides the grammar and instruments; the AI amanuensis provides continuous assistance and vigilance that lighten the cognitive and organisational burden of applying that grammar correctly over time. We first recapitulate AHWF's pillars and MWS as defined in earlier work, situating them in design-science research and innovation studies. We then characterise the AI amanuensis as a project-specific conversational partner whose capabilities in classification, memory, pattern detection, narrative reconstruction and meta-reflection are mapped, pillar by pillar, onto AHWF's demands. We describe the resulting AHWF+AI meta-artifact, detailing the assistance affordances that can reduce drift between intentions, operations and facts, preserve the layered distinction between processes and processings, reveal gaps in Selfs' Complexes, keep Value–Life complexes coherent, and maintain CI–CV handshakes as completable conversations. Stylised vignettes in technology transfer and regulatory sandboxes illustrate how this combined artefact works in practice. We propose an evaluation roadmap, building on the FEDS framework for design-science evaluation (Venable et al., 2016), and discussing implications and risks. The central claim is that AHWF, when paired

with a carefully configured AI amanuensis, can move the "from theories to tools" ambition one step further: not only are first principles turned into tools, but those tools themselves are supported by a conversational intelligence that helps practitioners apply them correctly, consistently and reflexively.

KEYWORDS

Design Science, Socio-Technical Systems, Completion, Accountability, Governance, AI Amanuensis

1. INTRODUCTION

We have long called for research that does not merely explain technological innovation and entrepreneurship, but that also produces tools practitioners can use (Romme and Holmström, 2023). This aspiration is closely aligned with the design science tradition in Information Systems, which emphasises the creation and evaluation of artefacts intended to extend human and organisational capabilities (Hevner et al., 2004). The Ad Hoc Weaving Framework (AHWF) was presented as one such meta-artifact (Bonatti *et al.,* 2025). It addresses a familiar "last-mile" problem: organisations generate rich theories, strategies and governance schemes, yet innovation initiatives falter because key commitments about what exists, what matters, who acts and how conversations complete remain implicit or fragmented across disciplines. AHWF responds by making those underpinnings explicit and usable. It offers a compact grammar for innovation practice built from five pillars and a Minimal Working Set (MWS) of three one-page instruments – the Trope Worksheet, the Value–Life Complex Canvas and the CI–CV Ladder – that teams can run in ninety minutes without heavy training.

The core promise of AHWF is that if innovation teams agree on what counts as an event, an operation, a relation, an intention, an activity, a fact, and a process or processing, and if they track the life-cycle and temporal posture of these items; if they explicitly separate natural processes from human processings; if they see persons and organisations as Selfs' Complexes with viable-system functions and faculties; if they keep mission, governance, learning, planning, doing and facts on one page; and if they design communication as sequences of completable conversations, then a surprising amount of ambiguity and slippage can be removed from innovation practice (Bonatti et al., 2025). The vignettes in the original AHWF paper – a biosensor transitioning from university lab to public hospitals, a corporate innovation portfolio, and a regulatory sandbox for privacy-preserving know-your-customer services – show how this grammar clarifies obligations, shortens cycles and improves gate decisions.

Yet, as the AHWF paper itself acknowledges, applying this grammar in live projects imposes a substantial cognitive and organisational burden. Someone must keep track, over months or years, of which tropes were defined, how their life-cycles and time-modes have evolved, whether intentions are still paired with actions and facts, whether Value–Life complexes remain aligned, whether CI–CV handshakes are being honoured, and where faculties and powers are missing. In practice, that burden falls unevenly, often onto a small number of reflective practitioners or external advisors, and gaps inevitably open.

In parallel, large language models (LLMs) have become able to operate as conversational project companions. They can read through long archives of documents, minutes and emails; classify statements; highlight contradictions and gaps; recover earlier decisions and rationales; and generate structured reflections. When configured as persistent assistants around a project, they begin to approximate the role of an AI amanuensis: a digital scribe and collaborator that "sits at the elbow" of human innovators, not as an oracle but as a memory steward, pattern spotter and meta-editor.

This paper asks what happens if we explicitly couple these two developments. What becomes possible if AHWF's grammar and instruments are treated as the "score" and the AI amanuensis as the "performer's assistant" – an instrument that helps humans apply the grammar correctly over time, under the messy conditions of real innovation work? More concretely, how can an AI amanuensis help innovation teams maintain trope classifications, life-cycle and time-mode status, VLC coherence, CI–CV completion and Selfs' Complex diagnostics, thereby reducing cognitive load and increasing the reliability of AHWF-guided work?

We answer these questions by articulating AHWF+AI as a composite meta-artifact. We first brief readers on AHWF's pillars and MWS, using the canonical definitions from Bonatti et al. (2025), and locate the framework in design-science and innovation literatures. We then characterise the AI amanuensis as a project-specific conversational partner and describe, pillar by pillar, which assistance affordances are possible and desirable. We show how these affordances can be instantiated in practice, and how they can help to lighten human researchers' and managers' burdens in "correct working" with AHWF. Stylised examples embedded in technology-transfer and sandbox contexts illustrate the configuration. Finally, we outline an evaluation plan and consider limitations and risks, including over-delegation, epistemic bias and the danger of confusing AI-supported bookkeeping with substantive decision-making.

2. THE AD HOC WEAVING FRAMEWORK: CANONICAL PILLARS AND INSTRUMENTS

The AHWF paper (Bonatti *et al.,* 2025) is already fully dedicated to explicating the framework; we therefore offer only a condensed recapitulation, focusing on those aspects relevant to AI assistance. Readers interested in philosophical detail and design-science positioning are referred to the original article.

AHWF was conceived as a meta-artifact that connects four foundational domains – ontology (what there is), axiology (what should matter), agency (who can decide and do) and communication pragmatics (how action completes and becomes fact) – and turns them into tools that fit on a single page. It does so through five pillars.

The first pillar is a trope-based meta-ontology with life-cycles and time-modes. Instead of talking about amorphous "things," AHWF asks teams to represent the situation in terms of seven trope families: events, operations, relations, intentions, activities/actions/acts, facts, and processes or processings, signaled respectively as e, o, r, i, a, f, p. Each family follows the same life-cycle pattern: what can be (the space of affordances), what should be (commitments or requirements), what is being done (instances unfolding in time), and what is done (closures). AHWF often names these stages with the morphology -abilia, -anda, the instance label and -ata, respectively. Every trope is also tagged with a time-mode: Given (Y) for inherited past,

On-Going (O) for present unfolding, Potential (Π) for possibilities not yet committed, and Prospective (Ω) for committed forward-oriented intentions.

The second pillar distinguishes natural processes from human processings and maps both onto a layered ontology. The framework separates the dynamics of the world from organised human doings, not as a dualism but as a pragmatic safeguard against category error. Processes and processings are distributed across logico-mathematical, physico-chemical, biological, physiological, psychological and layered psyche (egocentric, allocentric, cosmo-oriented). Relational maps and Angyal-style graphs help teams track how movements in one layer constrain and enable moves in another.

The third pillar models persons and organisations as Selfs' Complexes, layered entities with faculties and powers. Individuals and collectives share a template that includes substrates (physico-chemical, biological), mental and psychological faculties and psyche at the three scales. Organisational Selves are mapped explicitly to the Viable System Model's functions: operations, coordination, internal regulation, environment-sensing, and policy/identity (Beer, 1985). Faculties and powers – such as sensing, reasoning, deciding, allocating, coordinating and learning – are located across these layers, and their presence or absence becomes diagnostic for innovation viability.

The fourth pillar is the Value–Life Complex (VLC), an architecture that moves values from declaration to facts. VLC canvases have five blocks: finalities (mission, value commitments, requirements); governing and communicating (decision rights, oversight, cadence and communication infrastructure); observing, learning and knowing (signals, indicators, thresholds and evidence pathways); planning, resourcing, doing, delivering and innovating (concrete plans, allocations and improvement loops); and activity, action and fact (checkable closures). Cross-links ensure that signals feedback to mission, governance constrains and is informed by planning, and facts update learning and purpose.

The fifth pillar is a communication pragmatics labelled CI–CV, supported by an enabling engineering stack (ECHICO). Communication is treated not as the transmission of messages but as the design of completable conversations whose success is defined by acts done and facts registered. AHWF distinguishes primary and secondary conversations, handshakes with explicit acceptance conditions, and integration mechanisms. The ECHICO stack (Encoding, Channels, Handling, Interaction, Coordination and Outcome registration) ensures that no level – from technical transmission to organisational authority – is taken for granted.

These pillars are made concrete through a Minimal Working Set. The Trope Worksheet forces clarity on what is to be done, being done and done, and on time-modes; the VLC Canvas keeps mission, governance, learning, planning and facts on one plane; the CI–CV Ladder scripts the conversations and handshakes that must complete. Together they have been used in vignettes involving technology transfer, corporate portfolios and regulatory sandboxes, and they are accompanied by a FEDS-style evaluation plan with metrics such as closure latency, evidence coverage and rework ratio.

In summary, AHWF supplies a rich, portable grammar for innovation practice. What it demands, however, is that this grammar be applied consistently over time, across many small and large decisions. From a design science perspective, AHWF may be viewed as an artefact intended to support innovation practice through a structured grammar and associated instruments (Hevner et al., 2004). The remainder of this paper explores how an AI amanuensis can help with that demand.

3. THE AI AMANUENSIS: CAPABILITIES FOR SUPPORTING AHWF

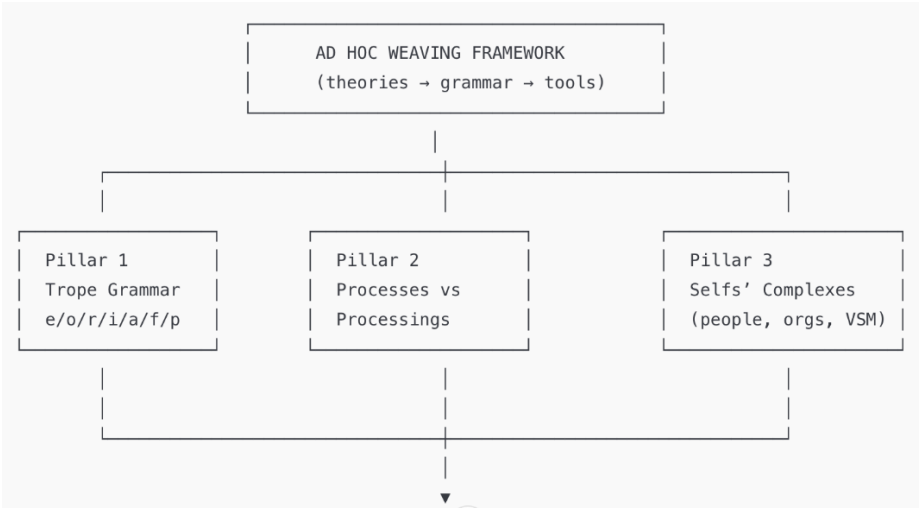
When we speak of an AI amanuensis, we mean a persistent, context-aware conversational agent based on a large language model that is configured around a particular innovation project or programme. It has access, under carefully governed conditions, to the project's textual artefacts – briefs, minutes, Trope Worksheets, VLC canvases, CI–CV ladders, reports, emails, code comments – and can engage with practitioners in natural language. Over time, it accumulates a memory of the project's conceptual weave and can act as a kind of externalised, structured memory and analytic engine.

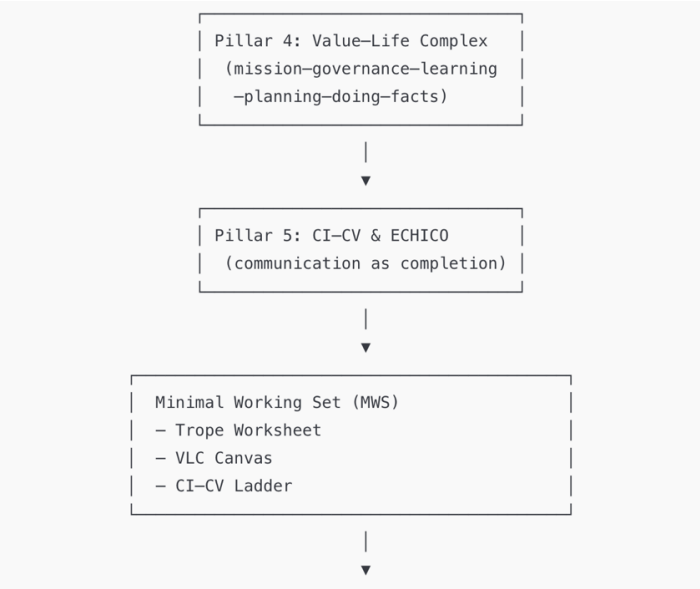
Within the constraints of contemporary LLMs, such an amanuensis can perform several tasks relevant to AHWF. It can read documents and map their components into AHWF constructs: identifying sentences that state intentions, operations, facts, relations or processes, and proposing their life-cycle and time-mode tags. It can trace where a value commitment first appeared, how it was later translated into an operation and whether a fact exists that would satisfy it. It can detect when the same term is used with different meanings in different places, or when two different terms refer to the same underlying trope. It can produce summaries and "weaving ledgers" of what has been decided, which obligations exist and which have completed.

These capabilities are not unique to AHWF; they are general properties of LLM-based assistants. What is distinctive here is that AHWF offers a target grammar toward which the AI's modelling can be oriented. Instead of simply clustering topics or extracting generic summaries, the amanuensis can be configured to ask specific AHWF-inspired questions: Is this clause an intention-anda or a fact-ata? Is this operation prospective or on-going? Does this conversational sequence culminate in a fact recognised by an appropriate authority? Is this concern located in the biological, psychological or cosmo-oriented layer? Which Selfs' Complex functions are being invoked or neglected?

The idea, then, is to define, for each AHWF pillar and instrument, a set of assistance affordances that an AI amanuensis could reasonably deliver. These affordances are not speculative; they are grounded in the tasks that AHWF expects humans to perform in Trope Worksheets, VLCs and CI–CV ladders. By making these affordances explicit, we can design more purposeful configurations of AI support and evaluate their effect on innovation practice.

4. PILLAR-BY-PILLAR ASSISTANCE: FROM TROPE GRAMMAR TO CI–CV





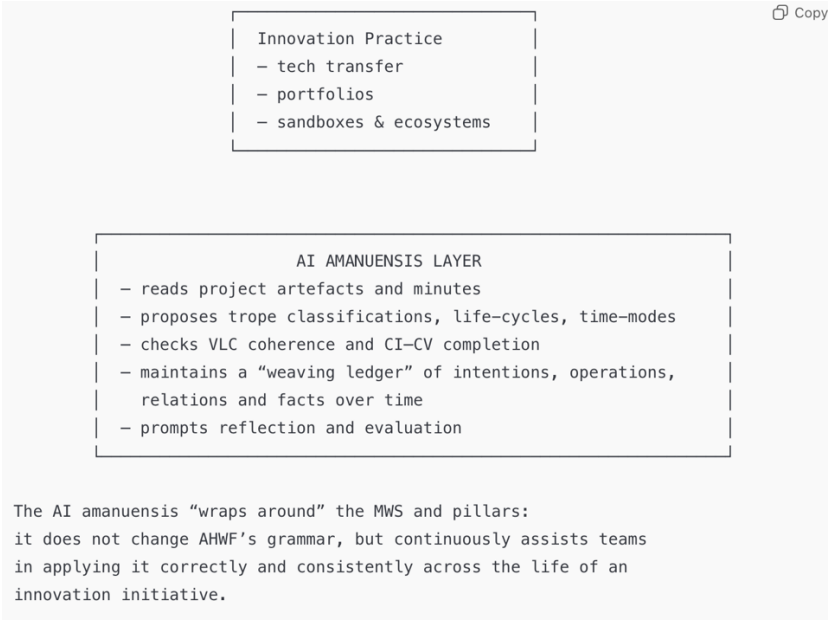


Figure 1. AHWF + AI amanuensis supporting innovation practice

4.1 Trope-Based Meta-Ontology: Classification, Status and Coherence

Applying AHWF's trope grammar in a live project involves a steady stream of micro-decisions. When someone writes "pilot with Hospital X completed", a colleague must decide whether this is a statement of an event, an operation, a relation or a fact; whether "completed" here signals a factual closure recognised by the hospital or merely an internal sentiment; and whether this pilot is Given, On-Going, Potential or Prospective at the time of the statement. When a project charter uses "commitment to fairness" in the mission paragraph and "parity tests" in the indicators block, someone must ensure that these items are connected through a chain of intentions, operations and facts.

An AI amanuensis can support this by performing continuous trope-aware reading and writing. When project teams draft Trope Worksheets, the AI can propose initial assignments of statements to e/o/r/i/a/f/p families and suggest life-cycle and time-mode tags, explaining its reasoning and inviting correction. When new documents are produced, it can scan for statements that resemble tropes and ask, "Should this be added to the Trope Worksheet; if so, how?" This reduces the friction of maintaining the worksheet as a living document.

Beyond classification, the AI can monitor life-cycle coherence. It can track whether intentions have corresponding operations and facts. If, several weeks into a project, the AI sees that "equitable access in public hospitals" remains only as an intention-anda, with no associated operations-anda or facts-ata, it can remind the team and suggest options, such as defining an operation for measuring price-to-income ratios and a fact that records public-hospital procurement under a threshold, as in the AHWF biosensor vignette. It can flag when an

operation is marked as "done" in narrative text but remains "being-done" in the worksheet, prompting reconciliation.

Similarly, the AI can guard time-mode discipline. It can detect when Potential items (Π) are being talked about as though they were Given, or when Prospective commitments (Ω) are made verbally but not recorded as such. In doing so, it lightens the burden on human facilitators, who otherwise must keep this entire classification in their head.

By keeping a continuous log of changes to trope statuses, the AI also creates a basis for later reflection. It can reconstruct, for a debrief, how many items moved from Potential to Prospective without ever being facted, or how often facts were claimed without evidence packets. This makes AHWF's life-cycle grammar not only a design tool but a basis for empirical analysis of project discipline.

4.2 Processes and Processings: Layered Reasoning and Category Vigilance

AHWF insists on distinguishing natural processes from human processings and on mapping both onto a layered ontology. In practice, teams often blur these distinctions. They may treat market shifts as if they were under governance, or regard organisational capacity constraints as immutable like physical laws. They may attribute a failure of coordination to "the system" rather than to a missing processing.

An AI amanuensis can help maintain category vigilance. When reading descriptions of risks, constraints or enablers, it can propose whether each is best treated as a process or a processing and at which layer. It can highlight where a natural process (such as sensor drift or model parameter instability) is being implicitly assumed to be controllable by policy, inviting explicit discussion of limits. Conversely, it can point out where human processings (such as governance procedures or training programmes) are being treated fatalistically.

In addition, the AI can generate multi-layer projection views of a situation. Given a Trope Worksheet and VLC Canvas, it can assemble a view that shows all biological constraints, all psychological considerations, all cosmo-oriented value commitments, and how they intersect. This helps practitioners see whether, for example, their design for a regulatory sandbox attends sufficiently to citizen comprehension (psychological and allocentric layers) as well as to cryptographic robustness (logico-mathematical and technical layers), as emphasised in the AHWF sandbox vignette.

Relational maps and Angyal graphs, which can be labour-intensive to draw by hand, can be generated semi-automatically by the AI from textual descriptions. Practitioners can then edit and annotate these diagrams rather than constructing them from scratch. This accelerates a pillar that is valuable but demanding.

4.3 Selfs' Complexes: Diagnosing and Supporting Viability

AHWF's third pillar models persons and organisations as Selfs' Complexes, aligned with Beer's Viable System Model. It asks teams to make explicit where System 1–5 functions reside, which faculties and powers are present, and where gaps exist. This is conceptually appealing but can be challenging in practice. Managers may resist mapping their organisation in this way; facilitators must translate conversations into System 1–5 structures and into lists of faculties and powers.

The AI amanuensis can assist by performing organisational self-analysis on textual material. It can scan role descriptions, organograms, minutes and governance documents, and propose a draft mapping of functions to units, suggesting, for instance, that a particular team is de facto performing both System 3 and System 4 functions, or that no person or committee appears to hold clear System 5 authority over a given innovation programme. It can propose lists of faculties and powers at stake in a project and highlight those that seem absent; for example, it might note that there are many references to sensing and deciding but few to learning and policy revision.

Because the variety of innovation projects is large, this analysis cannot be fully automated, but even a draft mapping can save substantial effort. In the technology-transfer vignette, for example, the AI could infer, from the description of the university, venture, hospital and regulator, which Selfs' Complex layers and viable-system functions each entity is likely to host, and which are missing from the joint governance structure. The project team can then confirm or correct this mapping, rather than creating it from scratch.

In reflection phases, the AI can reconstruct how the Selfs' Complex configuration changed over time – when a new System 4 function was created or when a System 2 coordination role was added – and correlate those changes with observed improvements in closure latency or evidence coverage. This links AHWF's organisational pillar with evaluation metrics suggested in the original paper.

4.4 Value–Life Complex: Maintaining Mission–Governance–Fact Coherence

The Value–Life Complex is arguably the most demanding of AHWF's pillars. It asks teams to write mission, governance, learning, planning and facts onto a single canvas and to maintain internal coherence as projects evolve. In the AHWF vignettes, the discipline of explicitly linking mission to indicators, to decision rights, to plans and to facts is what prevents equity and safety slogans from remaining mere rhetoric. But maintaining this discipline over time, as documents and decisions proliferate, is difficult.

An AI amanuensis can act as a VLC coherence engine. When a VLC Canvas is first created, the AI can help translate narrative mission statements into finalities expressed in trope grammar and can suggest candidate indicators that match those finalities. It can propose governance structures that align with viable-system logic, based on common patterns in similar cases. It can identify potential facts that would satisfy particular value commitments and propose evidence packets.

As the project progresses, the AI can monitor whether new decisions and activities remain consistent with the VLC. When minutes record a change in pricing strategy, the AI can ask whether this affects the equity indicators in the finalities block and whether those indicators and thresholds should be updated. When a new stakeholder is added, it can prompt an update to decision rights. When a pilot completes, it can ask whether the expected facts have been recorded and whether evidence meets the declared standard.

The AI can also help with cross-VLC analysis. In organisations running multiple innovation initiatives, each with its own VLC, the AI can identify common finalities, shared indicators and conflicting governance patterns, supporting portfolio-level weaving. This links to AHWF's suggestion that VLC logic can be applied at portfolio level, not just at initiative level, and that dynamic capabilities like sensing, seizing and transforming can be rendered observable in Value–Life terms (Teece et al., 1997; Bonatti et al., 2025).

4.5 CI–CV and ECHICO: Designing and Auditing Completion

The CI–CV pillar reframes communication as completion. It requires that conversations be designed so that obligations become recognisable completions with evidence. Primary and secondary conversations must be mapped; handshakes must be defined with acceptance criteria; technical and organisational means for encoding, handling, coordinating and registering outcomes must be specified. This is one of AHWF's most distinctive contributions, but also one of the most technically and organisationally challenging to apply.

An AI amanuensis can support CI–CV work in several ways. It can read descriptions of processes, meetings and interactions and suggest candidate handshakes, identifying where "updates" are being exchanged without clear acceptance conditions. For instance, in the sandbox vignette, the AI could recognise that model updates are described as being "notified to the regulator" and ask, "What is the acceptance condition for this notification; what evidence packet would constitute a completed handshake?" It can help teams script primary and secondary conversations, paraphrasing them into sequences of acts with clearly defined end-points.

In addition, the AI can monitor completion rates and failure modes. Because it has access to logs and documents, it can estimate closure latency for particular handshakes, count how many times operations are re-opened after supposed completion, and identify reasons, such as missing evidence. It can highlight conversations that produce many messages but few facts, prompting redesign. Over time, these analytics can feed directly into AHWF's evaluation metrics.

The ECHICO stack – Encoding, Channels, Handling, Interaction, Coordination, Outcome registration – provides a checklist for ensuring that technical and organisational conditions for completion exist. The AI can generate ECHICO templates for each handshake and check documents, systems and roles against them. For example, it can ask, "Through which channel will this consent be encoded, who handles it, how is interaction managed, who coordinates follow-up, and where is the outcome registered in a way that can be audited?" The burden of remembering each ECHICO element is thus shared between human facilitators and the AI.

5. THE MINIMAL WORKING SET WITH AN AI AMANUENSIS

The Trope Worksheet, VLC Canvas and CI–CV Ladder are where AHWF meets everyday practice. In the original paper, they are described in terms of human-run workshops and manual filling. We now consider how an AI amanuensis can be embedded in their use.

When teams run a Trope Worksheet session, they can treat the AI as a co-scribe. As participants speak, the AI records candidate tropes, suggests families, life-cycles and time-modes, and projects them on a shared screen for editing. This reduces the friction of transcribing and classifying in real time. After the workshop, the AI can tidy the worksheet, check consistency and generate a narrative summary that explains the status of each trope.

In a VLC Canvas session, the AI can be given a brief description of the innovation initiative and asked to generate an initial canvas, which the team then critiques and corrects. This transforms the task from filling a blank template to pruning, refining and contextualising a draft. As evidence is collected during the project, the AI can update the facts block and check that indicators and thresholds are still sensible.

For the CI–CV Ladder, which can be difficult to design from scratch, the AI can propose candidate primary and secondary conversations based on project artefacts and known patterns, such as stand-ups, hand-offs, approvals, audits and reviews. It can suggest handshakes and acceptance conditions and produce simple visual representations. The team then decides which to accept or modify. During execution, the AI can track actual communications against the ladder and highlight deviations.

In teaching and training settings, an AI amanuensis can scaffold learning. Students can be assigned a case, asked to run AHWF instruments, and allowed to consult the AI for hints and explanations, while still being required to make final judgements themselves. The AI's explanations can reinforce understanding of AHWF's grammar, not just produce answers.

6. STYLISED VIGNETTES: AHWF+AI IN ACTION

To illustrate the combined meta-artifact, we revisit the biosensor and sandbox vignettes from the AHWF paper and imagine them implemented with an AI amanuensis in the loop.

In the biosensor case, a university, a technology-transfer office (TTO), a new venture, a public hospital and a regulator must align around a path from promising bench results to deployment in public wards. With AHWF alone, the joint team uses the Trope Worksheet to distinguish processes such as sensor drift from processings such as trial execution, operations such as design-for-manufacture, relations such as licenses and partnerships, intentions such as equitable access and safety thresholds, and facts such as ethics approvals and purchase orders. They then build a VLC Canvas that ties finalities to governance, learning, planning and facts, and a CI–CV Ladder that scripts handshakes for design transfers, training and regulatory submissions.

With an AI amanuensis present, several aspects become smoother. When the Trope Workshop starts, team members speak freely; the AI listens, classifies and proposes entries in near real time. It notices that someone speaks of "partnership with Hospital X" in the present tense; it suggests that this is a potential relation (r-abilia, Π) rather than a given, because no contract or letter of intent has been mentioned. The team corrects or confirms this classification. Weeks later, when term sheets are drafted, the AI notes that the hospital is now engaged and suggests updating the relation to prospective (Ω) and then to fact-ata once the contract is signed and an evidence packet is assembled.

On the VLC canvas, the AI helps articulate equitable access in the finalities block as an intention-anda that will be satisfied only when a fact exists of a purchase order from a public hospital at a price-to-income ratio below a given threshold. When a pricing discussion arises, the AI reminds participants of this commitment and asks whether the proposed price would make that fact reachable. If not, the team must either adjust price, renegotiate mission or revise the indicator. The AI does not decide; it prevents implicit drift.

In CI–CV design, the AI proposes an ECHICO template for hospital training: encoding through teach-back tests, channels via in-person sessions and an online module, handling by named educators, interaction via role-play, coordination through the hospital's training office and outcome registration in a shared evidence repository. As training proceeds, the AI tracks attendance logs and performance rubrics and flags when teach-back performance fails to meet the acceptance condition. This allows the team to intervene before claiming completion.

In the sandbox case, firms and regulators use AHWF to articulate the sandbox charter in trope terms, build a VLC that makes public value auditable, and design CI–CV handshakes for consent, model updates and incident reporting. With an AI amanuensis, the regulator can ask for a machine-generated overview of all intentions-anda related to fairness and transparency, correlated with operations and facts. The AI may reveal that many firms mention transparency in mission statements but few have defined operations for making decision logic explainable to consumers. It can prompt the sandbox steering group to introduce new operations and facts, such as consumer-tested explanations and evidence of comprehension.

During the sandbox, firms submit model-update packets. The AI reads them and checks whether they contain the metrics and documentation specified in the handshake standard. If not, it alerts both firm and supervisor. It also monitors consent revocations and incident responses, computing closure latency and re-opening rates. At the end of each phase, it assists in producing the public report, ensuring that announced facts can be traced back to evidence packets.

These vignettes show that the AI amanuensis does not change AHWF's conceptual content. Instead, it helps to enact that content faithfully, by managing bookkeeping of meaning and enforcing the discipline of trope grammar, VLC coherence and CI–CV completion.

7. EVALUATION AND RESEARCH AGENDA

In design-science terms, AHWF+AI is an extension of the original meta-artifact. The evaluation of such artefacts is a central concern within design science research, where structured methodologies have been proposed for the design, demonstration and evaluation of innovative artefacts (Peffers et al., 2007). The evaluation of this extended artefact should follow the principles articulated in the AHWF paper, enriched with attention to AI-specific effects. The FEDS framework suggests combining artificial and naturalistic, formative and summative evaluations (Venable et al., 2016). Bonatti et al. (2025) already outline how to measure closure latency, rework ratio, evidence coverage and gate decision quality for AHWF alone. For AHWF+AI, the question is whether the presence of the AI amanuensis improves these metrics, and at what cost.

In artificial–formative settings, teams could be asked to apply AHWF to a case twice, once without and once with AI assistance. Independent judges would rate the clarity of trope classifications, the coherence of VLCs, the completeness of handshakes and the quality of rationales. Time taken and perceived cognitive load could also be measured. This would reveal whether the AI reduces effort and errors in applying the grammar.

In naturalistic–formative settings, pilot deployments of AHWF+AI in technology-transfer offices, innovation portfolios or sandboxes could be observed over several months. Researchers would track how often AI prompts led to corrections in trope status, mission–fact alignment or handshake design, and how practitioners experienced the assistance. Summative evaluation

would compare similar projects with and without AI support, measuring differences in closure latency, rework, evidence coverage and decision reversals.

Evaluation must also attend to qualitative phenomena. Does the AI amanuensis make practitioners more reflective about their own weaving, or does it lead them to defer too much to machine suggestions? Are certain voices amplified or marginalised by AI-mediated weaving? Does the discipline of AHWF, when combined with AI, risk imposing excessive formalisation, or does it liberate human attention for more substantive deliberation?

From a research standpoint, the AHWF+AI configuration opens several avenues. One is to study how conversational traces captured by the AI can be used to reconstruct weaving trajectories for scholarly analysis, extending sensemaking research (Weick, 1995; Cristofaro, 2022) into AI-mediated settings. Another is to investigate how AHWF and AI can be co-designed into digital platforms, such as Trustworthy Governable Platforms, so that conversational structures and AI assistance are integrated from the outset. A third is to explore variants of the AI amanuensis with different personalities, levels of assertiveness and pedagogical styles, and to test their effects on innovation outcomes.

8. RISKS AND ETHICAL CONSIDERATIONS

While we have emphasised the potential benefits of AHWF+AI, the configuration also carries risks. One is over-delegation. Practitioners might come to rely on the AI amanuensis for classification and coherence checks to such an extent that they stop engaging fully with AHWF's grammar. If the AI proposes that "transparency obligation X is an intention-anda at the allocentric layer" and the team accepts this without critical reflection, subtle biases in the AI's training or prompt design could shape the weave in unexamined ways.

Another risk is epistemic narrowing. The AI's suggestions are necessarily shaped by its training data and by the inputs provided. It may be better at recognising certain kinds of threads – for example, those that are well represented in technical and policy documents – and worse at others, such as tacit cultural practices or embodied experience. AHWF's layered ontology and Selfs' Complexes insist that psychological, allocentric and cosmo-oriented dimensions be considered; AI must be used in a way that does not crowd out human attentiveness to those dimensions.

A third concern is privacy and confidentiality. For the AI to function as an amanuensis, it must ingest sensitive project artefacts. This requires careful choices about where the model runs, how data is stored, who has access and how outputs are logged and audited. It also raises questions of authorship and accountability for AI-assisted outputs, especially when they inform regulatory decisions or public communication.

There is also a more subtle ethical issue of responsibility for meaning. AHWF emphasises that innovation teams must own their commitments, facts and handshakes. An AI amanuensis must therefore be designed and governed so that its role as a support tool is transparent, its limitations are acknowledged and its outputs are always subject to human review. The temptation to treat AI suggestions as authoritative must be resisted, especially when those suggestions concern value-laden decisions.

Addressing these risks requires governance at two levels: governance of the innovation programme (already in focus through VLC and CI–CV) and governance of the AI itself. The latter includes choices about model scope, logging, human-in-the-loop mandates and periodic audits of AI behaviour. In a sense, the AI amanuensis should itself be subject to an AHWF-style treatment: its purposes, processings, Selfs' Complex, Value–Life Complex and communication handshakes should be explicitly modelled, so that the tool remains accountable.

9. CONCLUSION

The Ad Hoc Weaving Framework was proposed as a way to "underpin the underpinning" of innovation practice by rendering ontological, axiological, agentive and communicational foundations explicit and usable. It offers a grammar and a Minimal Working Set that can be applied across domains, from lab-to-market transitions to regulatory sandboxes. Yet the very richness of this grammar means that applying it in live projects demands sustained attention and discipline.

The emergence of large language models as project-specific conversational partners provides an opportunity to share that burden. When configured as an AI amanuensis, such a model can help teams maintain trope classifications, life-cycle and time-mode status, layered distinctions between processes and processings, Selfs' Complex diagnostics, Value–Life coherence and CI–CV completion. AHWF tells us what to track; the AI amanuensis helps us actually track it.

This paper has sought to articulate this combined meta-artifact for a IADIS audience. We have recapitulated the canonical AHWF, characterised the AI amanuensis, mapped assistance affordances onto each AHWF pillar and instrument, and illustrated the result through vignettes. We have proposed an evaluation roadmap and reflected on risks. The intent is not to suggest that AI can or should replace human judgement, creativity or responsibility in innovation. Rather, it is to propose that, under carefully designed conditions, an AI amanuensis can make the demanding craft of AHWF-guided weaving more comfortable, more consistent and more accountable.

For IADIS' ambition to move from theories to tools, AHWF represented a first step: turning foundational theory into a usable grammar and instruments. AHWF+AI represents a second: equipping that grammar with a conversational intelligence that helps real teams use it correctly in the wild. The next step is empirical: to deploy this combined meta-artifact in diverse innovation settings, to measure its effects, to learn from its failures and to refine both the human and machine aspects of the weave.

REFERENCES

Beer, S. (1985). *Diagnosing the system for organizations*. Wiley.

Bonatti, M., Martin, M., & Jacucci, G. (2025). From theories to tools: The Ad Hoc Weaving Framework as a design-science meta-artifact for technological innovation, entrepreneurship, and technology management. *Proceedings of the OISI Workshop 2025*, Article 12. https://aisel.aisnet.org/oisiworkshop2025/12

Cristofaro, M. (2022). Organizational sensemaking: A systematic review and a co-evolutionary model. *Management Research Review, 45*(2), 215–238.

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly, 28*(1), 75–105.

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems, 24*(3), 45–77.

Romme, A. G. L., & Holmström, J. (2023). From theories to tools: Calling for research on technological innovation informed by design science. *Technovation, 121*, 102692.

Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal, 18*(7), 509–533.

Venable, J., Pries-Heje, J., & Baskerville, R. (2016). FEDS: A Framework for Evaluation in Design Science Research. *European Journal of Information Systems, 25*(1), 77–89.

Weick, K. E. (1995). *Sensemaking in organizations*. Sage.