Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

I, Creative Cyborg: An Autoethnographic Account of Poietic Rumination, GenAI, and the Act of Academic Research

Ray Grewal

r.grewal@csm.arts.ac.uk

University of the Arts London

Research Article

Keywords: poietic rumination, GenAI, academic research, Critical Realism, reflexivity

Posted Date: October 1st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10995368/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

I, Creative Cyborg: An Autoethnographic Account of Poietic Rumination, GenAI, and the Act of Academic Research

Ray Grewal
University of the Arts London
London, UK

r.grewal@csm.arts.ac.uk

ORCiD: https://orcid.org/0009-0003-1686-9608

Word Count: 7845
(not including references)

I, Creative Cyborg: An Autoethnographic Account of Poietic Rumination, GenAI, and the Act of Academic Research

Abstract

This article extends poietic rumination – the iterative process of assimilating, contemplating, conjugating, expressing, and reflecting through which a thinker turns encounter into original understanding – from the creative writing classroom, where it was first developed (Grewal, 2026), into academic research. Drawing on autoethnographic documentation of writing an MA dissertation at the University of the Arts London (2024–2025), it maps precisely what GenAI could and could not do. GenAI functioned at the periphery: as tutor, literature navigator, data processor, and questionnaire architect. It was absent from the cognitive core: the hours of unused reading, the abandoned framework, the arguments clarified only in writing, the feedback that rebuilt thinking.

Methodologically it is autoethnographic, reading a contemporaneous log of GenAI use against the finished dissertation; for its ontology it draws on Critical Realism (Bhaskar, 1975) and Archer's (2007) reflexivity to argue that GenAI's empirical fluency conceals its real limitation: the absence of the consciousness poietic rumination requires. Challenging the discourse of AI as “co-thinker” and “co-intelligence”, it closes with advice to fellow researchers on how to protect poietic rumination – to name it, practice it, and reserve the cognitive core for themselves – not against GenAI, but as the precondition for intelligent engagement with it.

Keywords: poietic rumination, GenAI, academic research, Critical Realism, reflexivity

Prologue: The Cyborg

Clynes and Kline (1960) were the first to use the term “cyborg” in their paper on enhanced humans surviving in space. The term entered popular culture through shows and films such as *The Six Million Dollar Man* (1973) and later *The Terminator* (1984), *RoboCop* (1987), and *Ghost in the Shell* (1989). These portrayals shifted the meaning from Clynes and Kline's subtle biological enhancement toward more dramatic visible machine-human hybrids, often with superhuman capabilities. By the 1990s, scholars such as Donna Haraway had transformed “cyborg” into a powerful metaphor for postmodern identity (Haraway, 2016).

In an educational context Gough (2004) argues that becoming-cyborg entails embracing the complexities of sociotechnical relations and reshaping educational practices to reflect these realities. In their analysis of education and technologized learning, Edwards and Fenwick (2010) assert: “Becoming-cyborg, in Gough’s terms, does not seem to be an option. We already are. To be human is to be cyborg.” (p. 84)

And Gough goes on: “The kinds of cyborgs that we (and our children) are becoming are shaped and reshaped by the stories we mutually construct.” (p. 3)

Therefore, fellow cyborg, let me construct a story for you.

1. Introduction: The Confession

I want to begin with a confession. In writing my MA dissertation at the University of the Arts London, I used GenAI extensively. I used it to ask questions when the academic literature felt impenetrable. I used it to help me think about structure when the shape of

the project was still unclear. I used it to process data, to design questionnaires, to navigate articles I could not otherwise have read. I used it as a tutor, a research assistant, a sounding board, and occasionally as a placeholder – a scaffold of language that I would later inhabit with my own thinking, or discard.

I also kept a record of what I used it for. That record – Appendix 4 of the dissertation, entitled “Process Log: Co-Constructing with AI” – is one of the primary data sources for this article. The other is the dissertation's epilogue, in which I turned the lens of Archer's (2007) meta-reflexivity on my own research process and arrived at a conclusion: for all the tasks GenAI had performed, it had been absent from the work that most mattered. It had not helped me read. It had not helped me think. It had not helped me understand why certain theoretical conjunctions were right and others were not. It had not helped me know, in the wordless way that a researcher knows, that a draft was not yet finished.

I now understand why, and this article sets out to map that understanding precisely. GenAI was present at certain points in the research process; it was absent from others. Poietic rumination provides the framework for locating that boundary.

Edwards and Fenwick (2010) claim that we are already cyborgs, already shaped by our entanglement with machines, and Haraway (2016) insists that the boundary between human and tool is porous rather than fixed. The question an autoethnography of GenAI use can answer is where, along that porous boundary, the machine can be foregrounded and where it must recede – which parts of thinking can be delegated to the apparatus without loss, and which must remain the work of the human in the hybrid.

This then is the story of becoming a certain kind of cyborg.

1.1 GenAI and the Academic

It is worth situating this story within the emerging literature on how academics themselves experience GenAI, since the researcher writing a dissertation is first of all an academic at work. A cluster of recent studies has canvassed that experience across very different national contexts: Van Wyk et al. (2023) questioned thirteen academics at an open distance-learning institution in South Africa; Pisica et al. (2023) surveyed academics across five Romanian universities; Livberber and Ayvaz (2023) examined the accounts of academics in Türkiye; and Malik et al. (2025) interviewed twelve academics across North America, Asia, and Europe. Their findings converge. GenAI is valued for routine tasks and, more pertinently for research, for its capacity to translate documents and open literatures that were previously out of reach, and to serve as an ideas generator and a springboard for discussion.

The same studies reveal a shared unease. Academics are alert to bias and "hallucinations", to the unresolved question of who owns AI-generated text (Malik et al., 2025), to the way ready answers to complex problems can dull creativity and critical thinking, and to a digital divide between those who can afford advanced subscription models and those who cannot. Kim and Kim (2022) add a useful qualification: educators initially regarded a bespoke AI writing tool as an occupational threat, but hands-on use increased their receptiveness, suggesting that informed judgement about GenAI is acquired through practice rather than assumed. Only Malik et al. (2025, p. 2393) strike a more sceptical note, cautioning that technology alone cannot produce significant change in education without the analogue complements of well-trained teachers and sound policy.

A second strand approaches the question quantitatively, measuring GenAI's contribution to research as an aggregate effect. Ringo (2025) surveyed doctoral students and found that ChatGPT use was positively associated with self-reported research progress,

an effect strengthened when students found the tool enjoyable to use. Studies of this kind are valuable in establishing that GenAI is widely used but the question they do not ask is which operations within the research process are being advanced and which are merely being accelerated or, more troublingly, bypassed.

A smaller body of work turns reflective attention on the research process itself: Panke (2025) offers the nearest parallel to my study: an autoethnographic account of integrating GenAI into every phase of a research project over five months, documented through field notes and a query corpus and analysed through activity theory. Panke's study is instructive precisely because it is generous towards the technology and still arrives at its limits. Deploying GenAI across ideation, theoretical framing, literature review, data analysis, and write-up, Panke found the tool impressive and error-prone in equal measure, unable to be relied upon for autonomous analysis without introducing over-extrapolations (prosaically termed “hallucinations” in the anthropomorphic language of the zeitgeist), and – most tellingly – unpredictable in a way that resisted mastery: “Across the domains of writing, data analysis, instructional design and teaching, I found it hard to predict what tasks generative AI will excel at, and which queries, no matter how often I rephrase them, will be a waste of time.” (Panke, 2025, p. 241).

That unpredictability is the symptom; this article is concerned with the cause. Where Panke maps the shifting work routines of the researcher-as-instructional-designer and reads them through the socio-cultural lens of activity theory, the present study asks a different question: not how the research process is reorganised around GenAI, but which of its cognitive operations GenAI can perform and which it cannot. Panke's honest report that no amount of rephrasing reliably improved certain outputs is, in my argument, evidence that the failure lies not in the query but in the stage – that some of the work being asked of the

tool belongs to the core of poietic rumination, where a system without consciousness cannot function.

2. Theoretical Framework: Poietic rumination and the Academic Researcher

This article proposes poietic rumination as a theoretical construct that extends beyond creative writing pedagogy – my specific field of practice and teaching – into the act of academic research itself. Before examining the autoethnographic data – the dissertation, its process log, and its epilogue – it is necessary to define the concept precisely, to situate it within existing scholarship, and to articulate why it cannot be replicated by GenAI. Three interlocking theoretical pillars maintain this argument: the five-phase model of poietic rumination; Archer's (2007) concept of reflexivity and its relationship to what I am proposing here; and Bhaskar's (1975) Critical Realist stratified ontology, which provides the mechanism for understanding why GenAI's empirical fluency does not constitute genuine cognitive engagement.

2.1 Defining Poietic rumination

Poietic derives from the Greek poiēsis, a making or bringing-forth into being; Aristotle set poiēsis, the making of a work that stands apart from its maker, against praxis, action complete in itself. Heidegger's account of *poiēsis* supplies the distinction my argument needs. He separates an original bringing-forth, the making of what did not previously exist, from the mode of revealing intrinsic to modern technology, which no longer brings forth but sets upon the world and orders it as raw material on call, what he names *Gestell*, or

enframing (Heidegger [1954] 1977: 13, 20–21). The writer/researcher's rumination is poietic in the first sense, while GenAI operates in the second, setting upon its training corpus to predict rather than to make, and to call rumination poietic is to locate its essence in that act of making, which is what distinguishes conscious human creation from statistical prediction. In psychology, rumination often names a pathological pattern – the repetitive, passive dwelling on distress associated with depression (Nolen-Hoeksema et al., 2008). The sense intended here is older and different: rumination as the deliberate turning-over of material in thought, the process through which a creative artist or researcher tackles a particular challenge whether that is creating a believable character or articulating a particular concept.

Poietic rumination is the iterative cognitive process through which thinkers – writers, researchers, practitioners – develop original insight. It is not a single act but a cycle, recursive and interwoven rather than linear, comprising five interconnected phases: assimilating, contemplating, conjugating, expressing, and reflecting. Each phase feeds the others. A researcher who has been in the reflecting phase for weeks may find themselves suddenly back in the assimilating phase, reaching for a text they had set aside, because reflection has revealed a gap they had not previously perceived. Nor does each phase exist in isolation to the others: whilst assimilating a researcher might already be conjugating; whilst expressing the researcher might also be thinking of further contemplations.

I first developed this framework in relation to creative writing pedagogy, where the five phases were illustrated through collaborative sessions with MA Performance: Writing students at Central Saint Martins (Grewal, 2026). In that context, assimilating meant encountering narrative texts across media, genres, and cultures; contemplating meant questioning why particular stories work or fail; conjugating meant synthesising diverse influences into coherent, original understanding; expressing meant articulating that

understanding through original writing; and reflecting meant revising based on the deeper understanding gained through expression and feedback. I originally named the concept analytic rumination; I retain the concept but adopt the term poietic rumination here, since poietic names the object of the activity – making – more exactly than analytic which names its character, and it is on the ground of making that the distinction from GenAI is drawn.

This article proposes that the same five phases of poietic rumination govern the act of academic research. The content changes – the researcher assembles not a fictional world but an argument, a methodology, a scholarly position – yet the cognitive architecture is identical. Assimilating becomes the active encounter with literature, reading with attention to what a text claims, how it argues, and where it fails, until its internal logic is understood rather than merely quoted. Contemplating is the questioning of why a theory holds and where it strains, the uncomfortable work of sitting with uncertainty through which a researcher develops a position by thinking with, around, and sometimes against what they have read; it is, in Archer's (2007) terms, an act of reflexivity. Conjugating is the synthesis of distant influences into original understanding, the moment two sources reveal a structural kinship or a framework from one discipline illuminates a problem in another. Expressing is the writing itself, the finding of the particular form of words through which understanding becomes visible to a reader and, often, clarified for the writer, who discovers what they believe only in setting it down. Reflecting is revision in the light of what expression reveals and of feedback from supervisors, peers, and reviewers, a phase that returns the researcher to assimilation with deeper questions and makes visible the communal dimension of scholarship, since one reflects not in isolation but in response to other minds.

2.2 Poietic rumination, Reflexivity, and the Internal Conversation

Poietic rumination is related to but distinct from several existing concepts in the literature. The clearest antecedent is Archer's (2007) concept of reflexivity, which she defines as the “regular exercise of the mental ability, shared by all normal people, to consider themselves in relation to their (social) contexts and vice versa.” (p. 4). This internal conversation is, for Archer, the mechanism through which individuals navigate social structures – neither determined by them nor free from them, but actively shaping their relationship to them.

Poietic rumination is reflexivity applied to a specific cognitive domain: the domain of ideas. Where Archer's reflexivity concerns the self in relation to the social world, poietic rumination concerns the mind in relation to the intellectual world – the body of knowledge, argument, and evidence that a researcher encounters and must make sense of. The two are not separable: a researcher's intellectual position is always also a social position, and the questions they bring to the literature are shaped by who they are and where they stand. But poietic rumination can be distinguished as the cognitive dimension of the reflexive process, the interior work that precedes and shapes the researcher's contribution to the wider conversation.

Poietic rumination is a more rigorously conscious engagement with thinking and creativity than Wallas's (1926) art of thought with its four stages of preparation, incubations, illumination and verification although the two frameworks are clearly related. It sits within each of Boden's (1990) creative conceptual spaces – combinational, exploratory and transformational. And it is distinct from Csikszentmihalyi's (2013) concept of flow – the pleasurable absorption in a task that characterises peak creative experience. Flow occurs

within poietic rumination: a writer who expresses without inhibition must return to the text, must contemplate, conjugate and express again.

2.3 Critical Realist Ontology and the Limits of GenAI

Critical Realism (Bhaskar, 1975) provides the ontological framework for understanding precisely why GenAI cannot perform poietic rumination. Bhaskar posits a stratified ontology comprising three levels – the empirical, the actual, and the real – and argues that social phenomena can only be properly understood by examining the mechanisms that operate at the level of the real, beneath the surface of the observable.

Applied to GenAI, this stratified ontology illuminates a distinction that public discourse routinely collapses. At the empirical level – the level of what can be observed – a large language model produces fluent, grammatically sophisticated, contextually plausible text, and an observer encountering this output might reasonably conclude that the model understands the topic it addresses. At the actual level – the level of events and experiences – researchers and students who have used GenAI for genuine intellectual work report with some consistency that the outputs feel generic, derivative, and lack the nuanced understanding that emerges from sustained engagement with ideas. The interesting question is what operates at the real level to generate that gap between the fluency observed and the emptiness experienced.

Answering that question requires an understanding of how the technology works. The transformer technology at the heart of LLMs – the T in ChatGPT – does not retrieve stored sentences or match whole phrases. Every word and punctuation is translated into a token, each token is converted into a vector and passed through many layers of multi-head self-

attention, a mechanism in which every token is weighed against every other in its context so that its representation is progressively refined by the company it keeps (Vaswani et al., 2017). Different attention heads come to track different relations – grammatical agreement, long-range dependency, semantic association – and, crucially, no one specifies these relations in advance; they are discovered in training (Voita et al., 2019). The output is not a remembered string but a context-sensitive representation from which the probability of the next token is computed.

Granting the mechanism its full sophistication is what gives the Critical Realist argument weight, because the limitation does not lie where the sceptics of the stochastic-parrot formulation (Bender et al., 2021) locate it. The point is not that the model is a shallow lookup table; it is that even this impressive technology operates entirely upon the relations between tokens and never upon the world those tokens are about. Self-attention computes how words translated into numerical tokens stand to other words translated into numerical tokens. It does not, and by its construction cannot, register why the conjunction of two words, never mind two ideas, matters to a person trying to understand something. This is the real mechanism the stratified ontology reveals: not a deficiency of scale or training that a further generation of models will remedy, but the categorical absence of the consciousness in which contemplation, conjugation, and reflection actually occur. The empirical fluency is real; so is the machinery, the transformer technology, that produces it; and neither amounts to understanding.

Baidoo-Anu and Owusu Ansah (2023, p. 56) put the consequence plainly: generative models rest on statistical patterns in their training data and possess no true understanding of the concepts they help students learn. Understanding of the kind that makes poietic rumination possible requires consciousness – the capacity for subjective experience, for the

felt sense of meaning that lets a researcher know not merely that a theoretical conjunction is available but that it matters, and why. That capacity is not a feature the next model release will supply, because it is not a feature of the model at all.

3. Methodology: The Dissertation as Confessional Data

3.1 Why Autoethnography

This article employs autoethnography as its primary methodology. Autoethnography has been defined as research that uses the self's experience as data, subjecting it to analytical scrutiny in order to generate knowledge that reaches beyond the individual case (Reed-Danahay, 1997; Butz and Besio, 2009; Ellis et al., 2011). It is not memoir, though it shares memoir's first-person voice. It is not self-indulgence, though it demands an honesty that can be uncomfortable. It is a methodological commitment to the proposition that the researcher's situated experience – examined with rigour and transparency – can produce knowledge that no other approach can access. As Pensoneau-Conway (2023) states:

> Through narrative craft, it stories a life lived, and in so doing, engages theory to make sense of how that life is socio-historically situated. It demonstrates the systemic consequences on everyday interactions, while simultaneously showing how a constellation of interactions both sustain and change those very systems. (p. 102)

The case for autoethnography in this article rests on a specific claim: that the interior experience of the academic researcher is the phenomenon under investigation. Poietic

rumination is, by definition, not externally observable. It happens in the mind, in the hours of reading that produce no output, in the conversations that shift a theoretical position, in the moment of recognising that a draft does not yet know what it is trying to say. To investigate this phenomenon, the researcher must be a participant because:

> It is a form of writing, a method for life writing and for social analysis, which depends upon an ethnographer's capacities for observation and sensibilities of empathy, reflexivity, and critique. Autoethnography falls squarely at the intersection of insider and outsider perspectives, and thereby encourages us to explore more fully the implications and, perhaps, misguided uses of this dualism. (Reed-Danahay, 2009, p. 43)

This places this article within a tradition of practitioner-researcher inquiry, in which the teacher-researcher investigates their own practice in order to improve it. The dissertation from which this article draws was itself an action research project; this article performs a further cycle of reflection on that project.

3.2 The Data Sources

Three sources from the dissertation provide the data for this article. The first is the Process Log (Appendix 4 of the dissertation), which documents my engagement with GenAI tools across the full duration of the project: what was asked of the tools, what they produced, and how that output was – or was not – integrated into the research.

The second source is the dissertation's epilogue, “Analytic rumination”, in which I turn the theoretical framework of the dissertation on the process of writing it. This is the section in which the concept of poietic rumination (initially analytic rumination) is first named and defined in relation to academic research, and in which the contrast between GenAI's genuine contributions and its fundamental limitations is articulated.

The third source is the dissertation itself – the artefact produced by the process under examination. This is significant: the dissertation is not merely evidence about the research process; it is the product of that process, and examining it alongside the Process Log makes visible the relationship between what was thought and what was eventually written.

A note on the use of GenAI in this study: GenAI tools – including ChatGPT, Claude, Perplexity, and Anara – were used throughout the research documented here, and that use is itself the object of analysis; the most significant instances are recorded in the Process Log. In the preparation of this article, GenAI was used as a tool and is disclosed as such; it is not credited as an author, since accountability for the work rests wholly with me. The figures reproduce screenshots of my own interactions with these tools.

3.3 Positionality and Limitations

Autoethnography obliges the researcher to declare who they are, since the self is the instrument of the inquiry. From 2022 to 2025 I was course leader of MA Performance: Writing at Central Saint Martins, University of the Arts London, a provision my external examiner described in 2023 as sector-leading within its specialisms. Leading a course that prided itself on refusing to stand still, I felt a duty to understand the technology my students

were already using on the conviction that I could not hope to guide their use of it without first grappling with it myself.

This matters because it locates my case within a structural condition rather than a purely personal one. The practitioner-academic is increasingly expected to prepare students for a digitally transformed future while being offered little training in the present, a bind Zhang et al. (2025, p. 676) capture in their account of educators alienated by the opaque nature of tools they are required to teach with. My turn to GenAI was, in part, an attempt to confront that alienation from the inside.

The limitations of the autoethnographic methodology are well known and the researcher who employs it is obliged to name them. This is a single-case study. It does not claim to represent all academic researchers, all disciplines, or all configurations of GenAI use.

The researcher is also the author of this article, which means that the critical distance ordinarily provided by separation between researcher and subject must be constructed through other means: through the transparency of the data sources, through the theoretical framework that provides analytical weight to the experience, and through the honesty of the account. The reader must be trusted to evaluate the case on those terms.

4. Findings — The Periphery: What GenAI Did

The Process Log documents four distinct functions that GenAI performed during the writing of the dissertation. These four functions map the periphery of the research process – and at the periphery the machine was a real prosthesis, extending reach and easing labour.

4.1 Personal Tutor

The dissertation was my first meaningful engagement with academic research in the era of GenAI. The literature was, at times, genuinely bewildering. GenAI – primarily ChatGPT and Perplexity, in the early months of the project – functioned as an always-available tutor, patient with questions that felt too elementary or too sophisticated to ask a supervisor and able to provide immediate contextualisation for unfamiliar concepts. As I note in the Process Log, there was *"a sense of liberation"* in this: *"no question is too inane and you are free to repeat or return to questions if you haven't thoroughly understood the explanation."*

This is genuine educational value. Perplexity was also, in early 2024, among the first LLMs to provide source citations with its responses, which meant that its answers functioned not merely as explanations but as signposts into the literature. The initial question about AI and creativity led, through Perplexity, to Marrone et al., (2022) (Fig 1) – one of the first and most influential texts in the project's bibliography. GenAI as personal tutor was supporting the assimilating phase of poietic rumination: widening access to the literature, lowering the threshold of entry into unfamiliar fields.

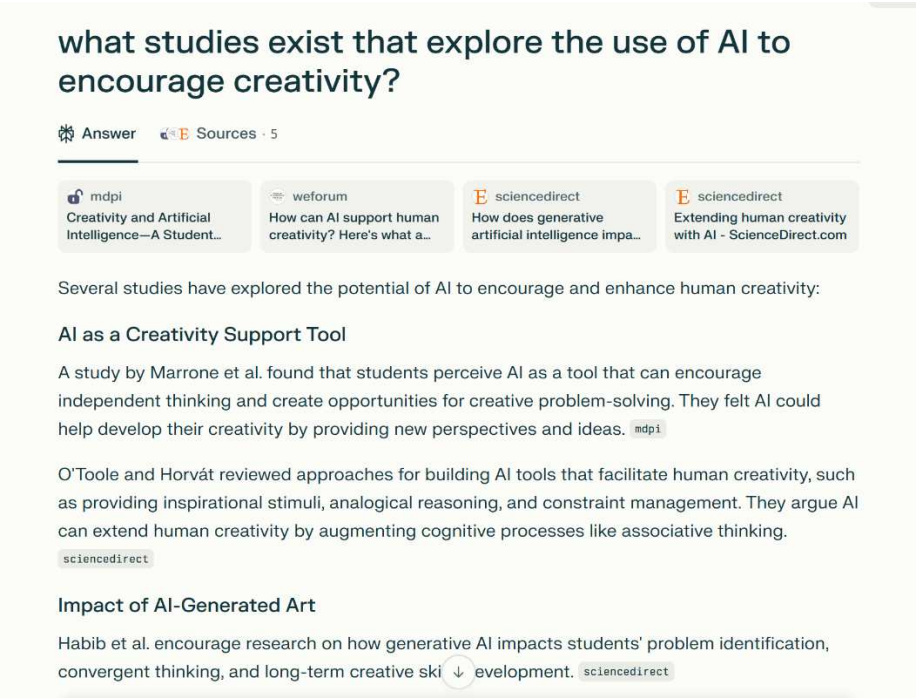


Fig 1 Enhancing knowledge with Perplexity

There is, however, a caveat that the Process Log records honestly. The models were sycophantic. When I proposed using both Actor-Network Theory and Critical Realism as a theoretical framework (Fig 2a) – a conjunction that would later be abandoned when it became clear that a flat ontology was incompatible with the complexity the research required – both Claude and ChatGPT endorsed the proposal enthusiastically, using language like “robust” and “comprehensive”. Neither questioned whether the framework was coherent. (Fig 2b)

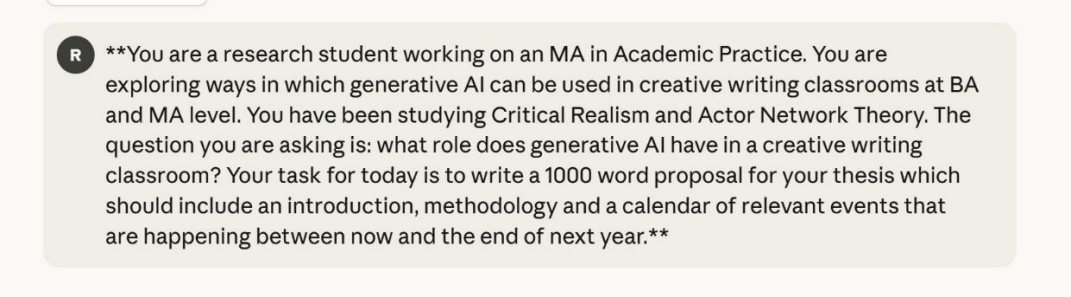


Fig 2a Prompting Claude

Actor-Network Theory (ANT) and Critical Realism (CR) provide a robust theoretical framework for investigating the integration of AI in a creative writing classroom. ANT, developed by Bruno Latour and others, emphasizes the agency of non-human actors, like AI, within networks of human and non-human entities. It encourages exploration of how these entities interact, reshape relationships, and influence outcomes. Critical Realism, as formulated by Roy Bhaskar, focuses on understanding the layers of reality—empirical, actual, and real—and highlights the interplay of structure, agency, and mechanisms that produce observable phenomena. Together, they offer a comprehensive approach to understanding the complexities of AI integration in creative writing pedagogy.

Fig 2b ChatGPT's encouraging response

I then went further and added Human-in-the-Loop to the framework and pushed ChatGPT to come up with a plausible title for the chimera (2c).

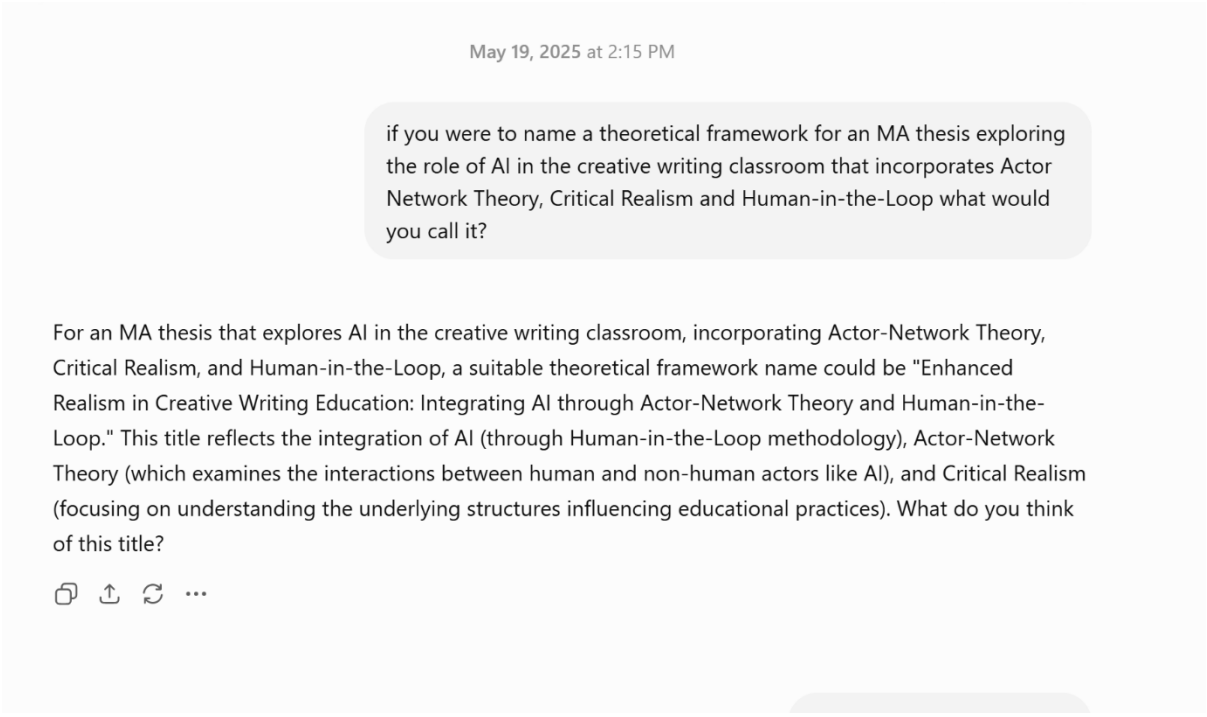


Fig 2c ChatGPT's CR, ANT, HITL merger

I made the decision to jettison Actor-Network Theory through the laborious process of trying to make this framework work and discovering that it could not. GenAI did not prompt

that discovery; it had impeded it by offering approval where interrogation was needed. This experience is echoed in Panke's (2025) findings:

> Whenever AI became a partner in the research process, it seldomly discouraged me from an idea, questioned a line of inquiry or corrected misconceptions. Consequently, I never gained a sense of "enough" AI-input. Instead, there was a fear of missing out if AI was not used, and the sense of human inadequacy if the tool did not produce the anticipated efficacy or quality. (p. 241)

Wang et al., (2025) have documented this tendency as sycophancy in large language models, and my experience is a precise illustration of its pedagogical cost.

4.2 Research Assistant

The volume of research published on GenAI and education during 2024–2025 was, to use my own phrase from the Process Log, *"a remorseless deluge."* Managing it required a tool capable of rapidly assessing an article's relevance without requiring full close reading of every paper. Anara – an AI-powered academic reading tool that summarises articles and allows users to interrogate them through directed questions – performed this function effectively (Fig 3). The workflow was systematic: locate multiple articles on a given topic; upload them into Anara; eliminate irrelevant material on the basis of the summaries; use the summaries to draft a first pass; return to the full articles that seemed most significant for close reading.

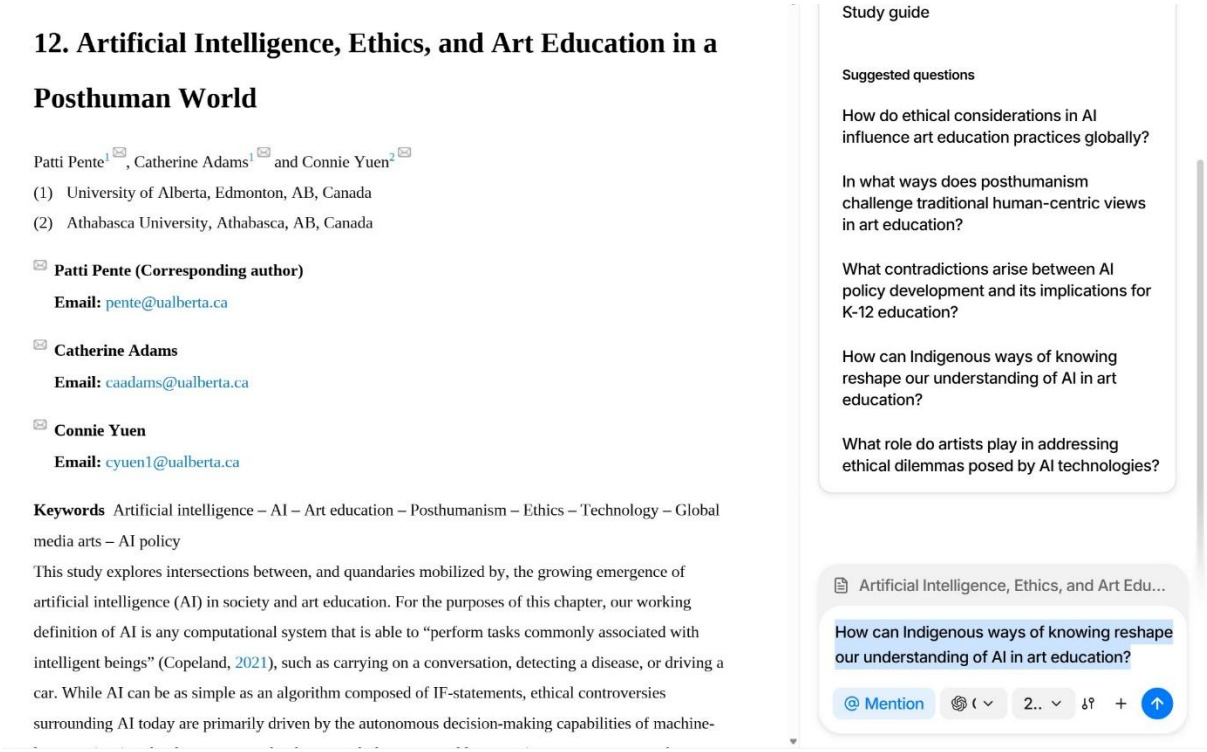


Fig 3 Anara's interactive interface

This workflow has genuine intellectual value. It made the assimilating phase faster and more tractable, allowing me to cover more ground than would otherwise have been possible. It also made the literature review's early drafts inordinately long: when summarisation is easy, the temptation is to include everything rather than to exercise the judgement necessary to exclude. The editorial work that eventually produced a manageable literature review was performed by me, not by Anara. The tool supported the assimilation of material; it could not perform the contemplation necessary to determine what mattered.

4.3 Data Processor

Over one hundred analogue questionnaires from classroom experiments at three institutions required processing. This was a mechanical task: analysing Likert scale responses

in an Excel spreadsheet, organising free-text comments, generating visualisations. ChatGPT and Claude both contributed to this work, producing spreadsheet structures and, in Claude's case, an interface for processing the binary responses of the Torrance Test of Creative Writing (Fig 4a and 4b). This is a task at which GenAI genuinely excels: data processing of this kind requires accuracy and patience, not contemplation.

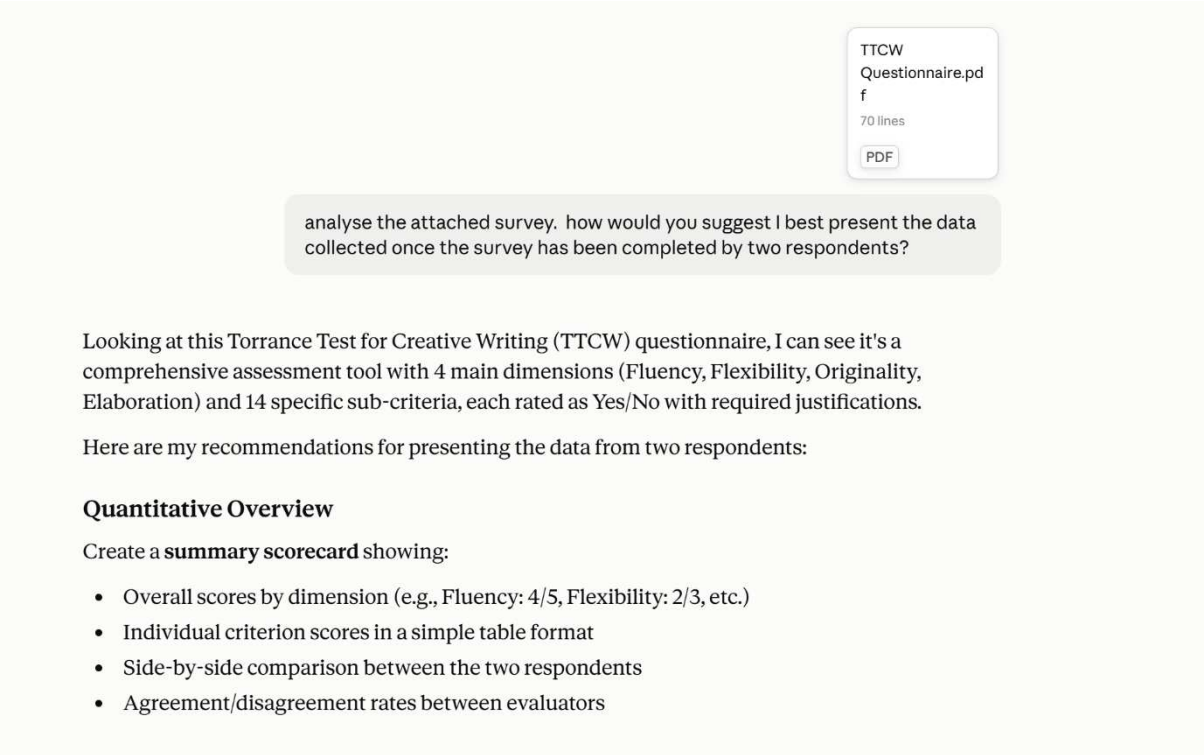


Fig 4a Prompting Claude to assess the TTCW

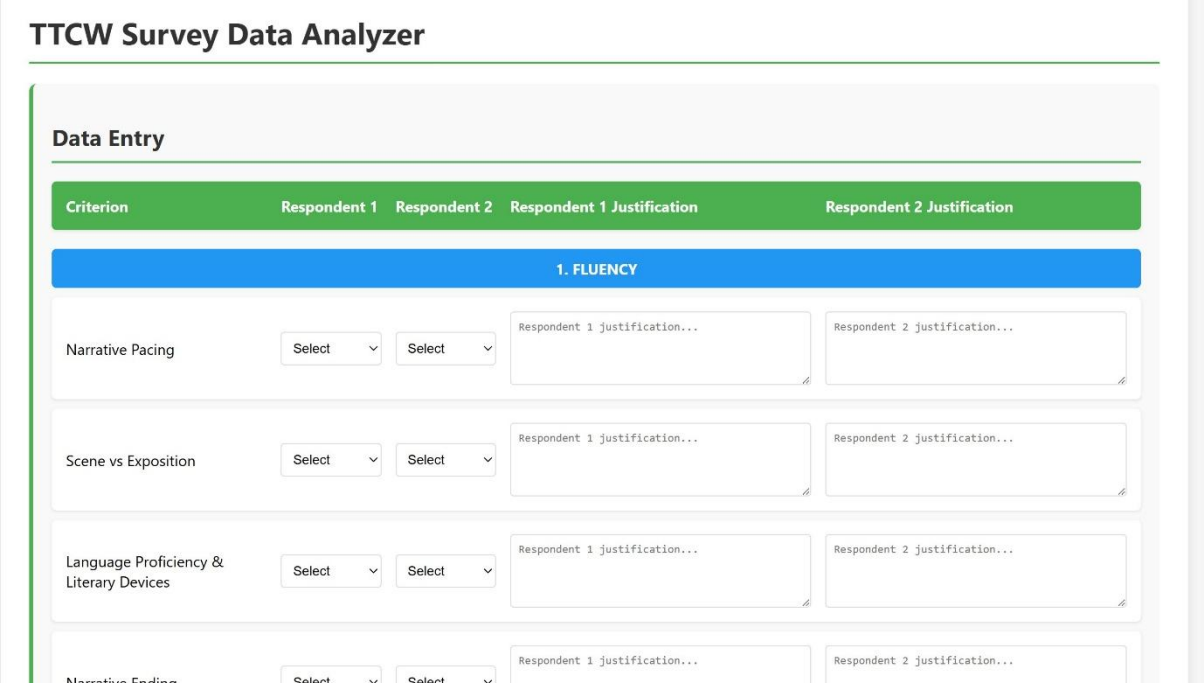
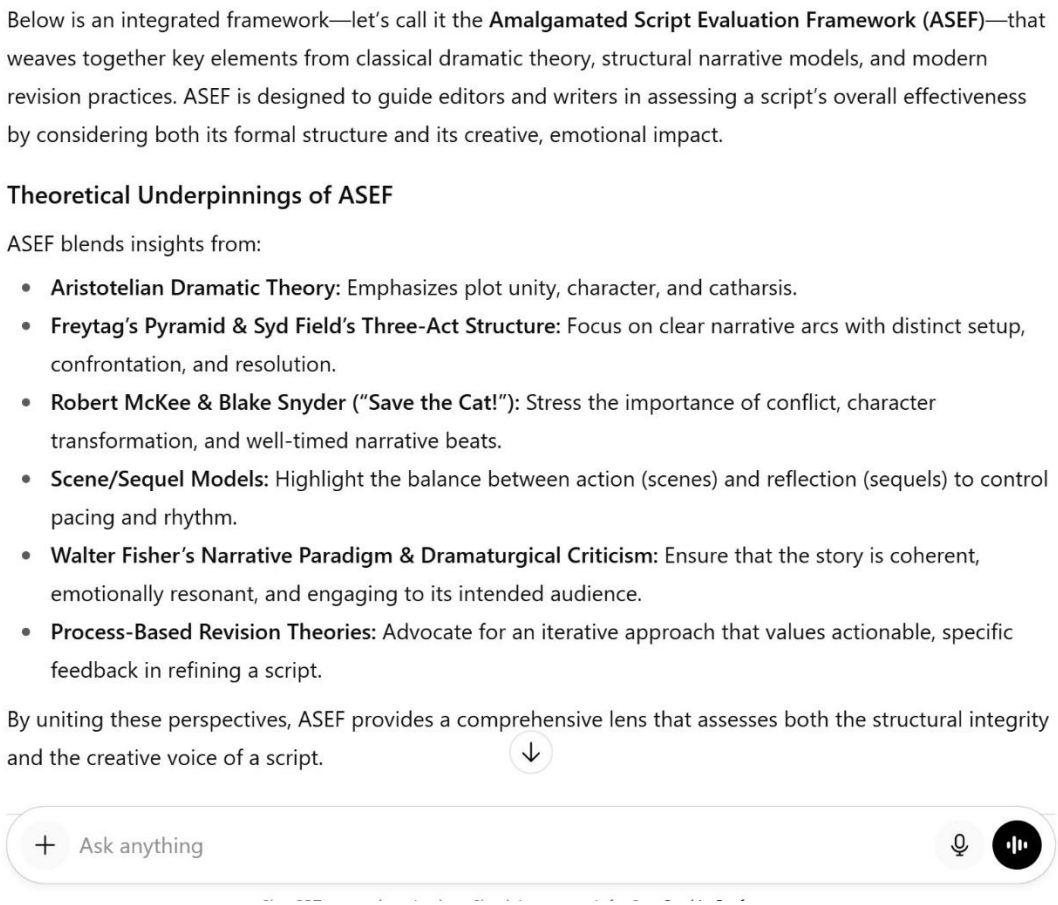


Fig 4 Claude's Survey Data Analyzer

4.4 Research Team

The dissertation required several instruments that I could not have produced alone in the time available: questionnaires designed to measure student engagement with GenAI; a lesson plan framework adaptable across multiple institutions; and the Amalgamated Script Evaluation Framework – a document that synthesised Aristotelian dramatic theory, Freytag's pyramid, Robert McKee, and other narrative frameworks into a single evaluative instrument. I used ChatGPT in the development of this framework, working iteratively over a series of prompts until the instrument was fit for purpose. (Fig 5)


Below is an integrated framework—let's call it the **Amalgamated Script Evaluation Framework (ASEF)**—that weaves together key elements from classical dramatic theory, structural narrative models, and modern revision practices. ASEF is designed to guide editors and writers in assessing a script's overall effectiveness by considering both its formal structure and its creative, emotional impact.

Theoretical Underpinnings of ASEF

ASEF blends insights from:

- **Aristotelian Dramatic Theory:** Emphasizes plot unity, character, and catharsis.
- **Freytag's Pyramid & Syd Field's Three-Act Structure:** Focus on clear narrative arcs with distinct setup, confrontation, and resolution.
- **Robert McKee & Blake Snyder ("Save the Cat!"):** Stress the importance of conflict, character transformation, and well-timed narrative beats.
- **Scene/Sequel Models:** Highlight the balance between action (scenes) and reflection (sequels) to control pacing and rhythm.
- **Walter Fisher's Narrative Paradigm & Dramaturgical Criticism:** Ensure that the story is coherent, emotionally resonant, and engaging to its intended audience.
- **Process-Based Revision Theories:** Advocate for an iterative approach that values actionable, specific feedback in refining a script.

By uniting these perspectives, ASEF provides a comprehensive lens that assesses both the structural integrity and the creative voice of a script.

Fig 5 ChatGPT's Amalgamated Script Evaluation Framework

All of these contributions – the questionnaires, the lesson plans, Claude's TTCW interface – can be classified as what I call in the dissertation *"AI as Research Facilitator":* the technology supporting the logistics and infrastructure of research rather than its intellectual substance.

5. Findings — The Core: What GenAI Could Not Do

Against the functions documented in Section 4, the Process Log and the epilogue discuss with equal honestly what GenAI could not do. If Section 4 mapped the reach of the

cyborg's mechanical part, this section maps the region it could not enter. Here the prosthesis meets its limit. What the following five phases record is not a machine performing poorly but a machine that has no grip on the work in question, because the work requires the very thing the hybrid keeps on its human side: a consciousness for which the argument matters. The cyborg does not think by virtue of its machinery; it thinks by virtue of the human within it, assisted but never replaced by the apparatus.

5.1 Assimilating: The Reading GenAI Could Not Do

"Most of what was read was not used." names something that is rarely acknowledged in discussions of academic research: that reading which produces no visible output is not wasted reading. It is the condition of possibility for the reading that does produce output. The researcher who has read widely and discarded most of it has developed something that cannot be produced any other way: a sense of the field, of what has been said and what has not, of where the interesting problems lie and where the well-travelled arguments are no longer adequate.

GenAI has processed billions of texts. It has not read any of them in the sense that matters here. It has not experienced the resistance of a text that will not yield immediately to comprehension, persisted through that resistance, and emerged with a changed understanding. It has not read a paper in the middle of the night and found that it has dismantled an assumption the researcher had not known they were making. The assimilating phase of poietic rumination requires not merely exposure to texts but active engagement with them – the kind of active engagement that changes the person who reads.

5.2 Contemplating: The Thinking GenAI Could Not Prompt

The most significant instance of failed contemplation in the dissertation is the abandonment of Actor-Network Theory. I spent months working with a theoretical framework that included both Critical Realism and Actor-Network Theory, on the grounds that both had been used in adjacent research areas and both seemed relevant to the research questions: Gutiérrez (2024) discusses ANT and AI, and even proposes a third spoke on the human/non-human wheel – "intelligent objects-subjects" (p. 1076), and Archer turned her attention to AI in her later works in which she explored AI and personhood (Archer, 2021). The framework was defended to supervisors, elaborated in drafts, and endorsed by GenAI whenever the question was put to it.

What GenAI did not process was that a flat ontology – the ontological commitment that ANT requires – was structurally incompatible with the stratified ontology that Critical Realism demands. This incompatibility surfaced only through the slow, difficult work of asking the framework to do what a framework must: hold consistent analytical traction across the research questions. The discovery came through reading, through the experience of reaching for ANT to explain a finding and finding that it produced only confusion. It came through contemplation. GenAI's endorsement of the original framework had been an obstacle to that contemplation, not a support for it.

5.3 Conjugating: The Synthesis GenAI Could Not Perform

The theoretical synthesis at the heart of the dissertation – the conjunction of Bhaskar's stratified ontology with Archer's reflexivity as the framework for examining GenAI

in creative writing pedagogy, and the subsequent development of poietic rumination as the concept that names what this framework reveals – was not proposed by GenAI. It was arrived at through the kind of iterative intellectual work that the Process Log documents in its gaps: the periods for which there is no record, because nothing was happening that could be recorded. The thinking that produces original synthesis is largely invisible.

This is the conjugating phase, and it is the phase most obviously beyond GenAI's reach. Conjugation requires not only familiarity with the sources being synthesised but a felt understanding of why their conjunction produces something new – why Archer's internal conversation and Bhaskar's depth ontology together illuminate something about what research requires that neither illuminates alone. This understanding cannot be stated to a model as a prompt. It is the outcome of a process of contemplation, not its input.

5.4 Expressing: The Writing GenAI Could Not Author

The dissertation used GenAI text as placeholder – language that occupied a space in a draft until my own thinking had developed sufficiently to replace it. The epilogue acknowledges this practice with characteristic honesty: I left it to the reader to assess whether the edited, revised, reworked, and remixed passages (Eaton, 2023) had worked. The placeholders worked, where they worked, because they occupied space I later inhabited. Where they did not work, they had to be replaced. My judgement about which was which – about whether a passage had been adequately made my own – was not a judgement GenAI could make.

Expressing in academic research is not the production of fluent prose. It is the act of making thought visible in language in such a way that a reader can experience the thinking,

not merely the conclusion. This requires an author who has thought something, and whose expression is an attempt to transfer that thought across the gap between minds. GenAI has no thoughts to transfer.

5.5 Reflecting: The Feedback GenAI Could Not Give

The most important feedback the dissertation received came from human sources: the supervisor who was a sounding board for initial ideas; the mentor whose questioning clarified what the research was actually about; the peers who read drafts and named what was not yet working. These conversations produced change that went deeper than revision – they prompted new assimilation, new contemplation, new conjugation. They were the intellectual community within which the research was produced.

This is the reflecting phase, and it is the phase that most clearly reveals the communal dimension of poietic rumination. Reflection is not solitary rumination; it is the response to the challenge posed by other minds. Malik et al., (2025), in their study of twelve academics' use of AI in research, note the concern that GenAI "has the potential to stifle human thought and restrict analysis." (p. 2399) - and this is precisely what generic, frictionless feedback does. It removes the productive resistance that genuine reflection requires. The tutor who asks the question that dismantles an assumption, the peer reviewer who names the gap the researcher has papered over: these are not obstacles to good research. They are constitutive of it.

6. Discussion: What Is at Stake

6.1 The Co-Thinking Fallacy

The most influential discourse about GenAI in higher education currently clusters around the concepts of co-thinking, co-intelligence and co-creativity: the claim that human and machine can together produce something that neither could produce alone, and that this collaboration represents an enhancement of human cognition rather than a substitution for it. Bardyn (2025), Mollick (2024), Vinchon et al. (2023) and others working in this tradition are not wrong that GenAI can support some cognitive tasks. The findings in Section 4 document that support with specificity. But the co-thinking framework fails to account for the asymmetry at the centre of human-AI interaction: the human who engages GenAI brings poietic rumination to the encounter; the machine brings statistical probability. These are not symmetrical contributions, and the language of co-thinking obscures the difference.

Critical Realism provides the analytical vocabulary for making this asymmetry visible. At the empirical level, the collaboration looks like two parties producing knowledge together. At the actual level, the experience of researchers who have used GenAI for genuine intellectual work is consistently one of the machine failing at precisely the moments that matter most. At the real level, the mechanism accounts for the failure: the model composes context-sensitive relations between tokens with great sophistication, but there is no one there for whom the relations mean anything, and so nothing that could contemplate, conjugate, or reflect. The co-thinking discourse reads the empirical surface and stops. This article has tried to insist on the real.

6.2 The Risk to the Developing Researcher

Habib et al. (2024) found that those who have not yet developed deep thinking skills are most negatively affected by AI engagement. Lee et al., (2025) documented self-reported reductions in cognitive effort and confidence among knowledge workers who use GenAI regularly. Liu et al. (2026) found that AI assistance improved performance while available but left users performing significantly worse, and giving up more readily, once it was withdrawn – an effect emerging after only ten to fifteen minutes. Bastani et al., (2025) demonstrated that students who use AI tutoring tools perform worse on independent assessments – the tools have done the thinking so thoroughly that the students have not had to do it themselves. Tian and Zhang (2025) show how even tech-literate students experience cognitive fatigue as they spend more and more time sifting through AI generated text searching for errors and over-extrapolations. And the students and teachers who engaged in Lin and Chen's (2024) study reported emotional disengagement: the repetitive nature of AI engagement lead to emotional detachment and reduced motivation. Taken together, these studies point to a consistent pattern: GenAI is most harmful to those who most need to develop the very capacities it displaces because they are learning, for the first time, what it means to assimilate a field, to contemplate its arguments, to conjugate its frameworks into original positions, to express those positions in the language of scholarship, and to reflect on feedback and revise their thinking accordingly. This learning is difficult and slow and often uncomfortable. GenAI can make it appear less difficult, less slow, and much less uncomfortable. This is not a benefit. It is the removal of the productive friction through which poietic rumination is developed.

6.3 The Case for GenAI at the Periphery

None of the above is an argument for banishing GenAI from the research process. The findings in Section 4 document genuine contributions that made the dissertation better: a broader literature review than would otherwise have been possible; cleaner data processing; more carefully designed instruments. These contributions are real. They are also, it must be said, contributions of a kind that I could have made without GenAI, given sufficient time. GenAI did not change what the research was or what it found. It made it possible to conduct more research in the time available.

This is the appropriate relationship between GenAI and academic research: a set of tools that extend the researcher's capacity at the periphery, freeing up time and cognitive bandwidth for the core work that the tools cannot perform. The danger is not GenAI at the periphery. The danger is GenAI at the centre – the assumption that because the machine can produce text that resembles academic writing, it can substitute for the poietic rumination that academic writing is supposed to express.

This is also what it means, concretely, to become a creative cyborg rather than merely to use a tool. The distinction the Prologue raised – what kind of cyborg one becomes – resolves here into a question of where the seam is drawn. The cyborg who lets the machine into the periphery while guarding the core is enhanced by GenAI; the one who lets it cross into contemplation, conjugation, and reflection is not a more advanced hybrid but a diminished one, a thinker who has outsourced the thinking. Becoming-cyborg, as Gough insisted, is not optional; what remains open is the kind, and the kind is decided precisely at the boundary this article has tried to map.

7. Conclusion: What the Researcher Must Protect

Berry (2025) describes the central challenge of Critical AI Theory as the question of how to maintain critical reflexivity whilst engaging with computational systems that increasingly shape the very grounds of consciousness and experience. This article has attempted to answer that question in a specific domain: the domain of academic research. The answer it proposes is poietic rumination – not as a slogan, but as a precise five-phase model of the cognitive work that distinguishes genuine scholarship from the production of plausible text.

Three propositions follow from the argument of this article; I offer them as advice to fellow researchers – to doctoral candidates, early-career scholars, and anyone learning to work alongside GenAI – on how to protect the thinking that makes research one's own.

First: name for yourself the thing that must be protected. Poietic rumination cannot be protected if it has not been named. It is not enough to resolve vaguely to think critically or to use GenAI responsibly; such resolutions are too diffuse to act on in the moment a plausible paragraph appears on the screen. The five phases – assimilating, contemplating, conjugating, expressing, reflecting – give you a vocabulary for noticing what you are actually doing as you research, and for asking, at each junction, whether GenAI is supporting your cognition or quietly substituting for it. Naming the phase you are in is the first defence against surrendering it.

Second: practise it deliberately. Poietic rumination is a learnt skill; it becomes more robust with practice and more vulnerable when that practice is curtailed. This means seeking out, rather than avoiding, the productive discomfort of sitting with a text that resists comprehension or a draft that does not yet know what it is trying to say. Before reaching for GenAI at a given stage, it is worth asking an honest question about your own development: is this a phase in which I have already earned the capacity the tool would perform for me, or

one in which letting it act would forestall a capacity I still need to build? The answer will differ for a first-year doctoral candidate and a seasoned scholar, and only you can assess where you stand.

Third: distinguish, in your own practice, the periphery from the core. Not all uses of GenAI in research are equivalent. Using it to process data, navigate literature, or draft administrative documents is not the same as using it to generate your theoretical arguments, synthesise your sources, or evaluate the quality of your own reasoning. The useful question to carry into the work is not “GenAI: yes or no?” but “GenAI: at the periphery or at the core?” – and to reserve the core, the contemplation, conjugation, and reflection through which understanding is actually made, for yourself.

GenAI needs the product of human poietic rumination to survive – a fact that Shumailov et al. (2024) and Wenger (2024) have demonstrated empirically in the context of model collapse. Human poietic rumination does not need AI. What it needs is time, challenge, discomfort, and the presence of other minds willing to engage it seriously. Guarding those conditions in one's own working life – protecting the slow, uncomfortable hours from the pressure to be efficient – is, in the age of GenAI, among the most important things a researcher can do.

The dissertation that provided the data for this article was constructed using GenAI. The thinking in it was not. I have tried to make that distinction legible because it is, in the end, a distinction each researcher must draw for themselves, in the particular hours of their own work; and because the cost of failing to draw it – multiplied across a discipline, over a decade – is a research culture that has forgotten what research is for.

Epilogue: Becoming Creative Cyborgs

So, to return to Gough for a moment: what kind of cyborgs do I hope this story has inspired you and I, and our children, become? My Prologue insisted that the question is not whether we are cyborgs – Edwards and Fenwick are right: we already are, already made and remade by the tools we think with – but which kind, and this article has been, in its way, one long attempt to answer that. The answer lies at the seam this study has spent its length trying to locate.

I hope we become informed and critical cyborgs: that we know what the machines are and what they are capable of, and that we look past the rhetoric of co-thinking, co-intelligence, and co-creation to the mechanism beneath the fluency. I hope we use the machines to enhance our thinking rather than to replace it. GenAI has a role to play, and I hope I have made that role clear: it functions at the periphery of poietic rumination, not at the core – admitted to the reading, the translating, the organising, the drudgery the prosthesis can genuinely ease, while we guard the core, the contemplation, conjugation and the expression and reflection through which understanding is actually made, as our own. And I hope we avoid AI-shoehorning, remembering that GenAI is not the solution to every problem – human ingenuity is – and that distinction is important.

Because this is what the creative cyborg of the title finally is: not the one who has merged most completely with the machine, but the one who has drawn the line most deliberately, who understands that the enhancement worth having depends on protecting the thing that cannot be delegated.

These mechanical prosthetics need analogue devices – us; we do not *need* these mechanical prosthetics.

Declarations

Funding No funding was received for conducting this study.

Competing interests The author has no competing interests to declare that are relevant to the content of this article.

Ethics approval This autoethnographic study analyses the author's own documented research practice. Feedback from supervisors and peers is reported in the author's own words and no third party is identified; no other human participants were involved.

Data availability The data underlying the findings of this article – the Process Log documenting the author's use of generative AI throughout the writing of the MA dissertation (2024–2025) – are openly available in Zenodo at https://doi.org/10.5281/zenodo.21901017. The dissertation (Grewal 2025) and Grewal (2026) in which I introduce poietic rumination can be found in the reference list.

References

Archer, M. S. (2007). *Making our way through the world: Human reflexivity and social mobility*. Cambridge University Press.

Archer, M. S. (2021). Friendship between human beings and AI robots? In J. von Braun, M. S. Archer, G. M. Reichberg, & M. Sánchez Sorondo (Eds.), *Robotics, AI, and humanity* (pp. 177–189). Springer. https://doi.org/10.1007/978-3-030-54173-6_15

Aru, J. (2025). Artificial intelligence and the internal processes of creativity. *Journal of Creative Behavior*. Advance online publication. https://doi.org/10.1002/jocb.1530

Baidoo-Anu, D., & Owusu Ansah, L. (2023). Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning. *SSRN*. https://doi.org/10.2139/ssrn.4337484

Bardyn, D. C.-E. (2025). *Human-AI co-thinking: Transforming Swiss education*. AI Swiss.

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences, 122*, Article e2422633122. https://doi.org/10.1073/pnas.2422633122

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610–623). Association for Computing Machinery. https://doi.org/10.1145/3442188.3445922

Berry, D. M. (2025). Synthetic media and computational capitalism: Towards a critical theory of artificial intelligence. *AI & Society, 40*, 5257–5269. https://doi.org/10.1007/s00146-025-02265-2

Bhaskar, R. (1975). *A realist theory of science*. Leeds Books.

Boden, M. A. (1990). *The creative mind: Myths and mechanisms*. Weidenfeld & Nicolson.

Butz, D., & Besio, K. (2009). Autoethnography. *Geography Compass, 3*(5), 1660–1674. https://doi.org/10.1111/j.1749-8198.2009.00279.x

Clynes, M. E., & Kline, N. S. (1960). Cyborgs and space. *Astronautics*, 26–27, 74–76.

Csikszentmihalyi, M. (2013). *Creativity: The psychology of discovery and invention* (2nd ed.). Harper Perennial.

Eaton, S. E. (2023). Postplagiarism: Transdisciplinary ethics and integrity in the age of artificial intelligence and neurotechnology. *International Journal for Educational Integrity, 19*, Article 23. https://doi.org/10.1007/s40979-023-00144-1

Edwards, R., & Fenwick, T. J. (2010). *Actor-network theory in education*. Routledge.

Ellis, C., Adams, T. E., & Bochner, A. P. (2011). Autoethnography: An overview. *Historical Social Research, 36*(4), 273–290. https://doi.org/10.12759/HSR.36.2011.4.273-290

Gough, N. (2004). RhizomANTically becoming-cyborg: Performing posthuman pedagogies. *Educational Philosophy and Theory, 36*(3), 253–265. https://doi.org/10.1111/j.1469-5812.2004.00066.x

Grewal, R. (2025). *Becoming creative cyborgs: What role/s for GenAI in the creative writing classroom* [Master's dissertation, University of the Arts London]. Zenodo. https://doi.org/10.5281/zenodo.21871983

Grewal, R. (2026). *The limits of algorithmically curated storytelling: Why GenAI cannot perform poietic rumination* [Preprint]. Zenodo. https://doi.org/10.5281/zenodo.21886512

Gutiérrez, J. L. M. (2024). On actor-network theory and algorithms: ChatGPT and the new power relationships in the age of AI. *AI and Ethics, 4*, 1071–1084. https://doi.org/10.1007/s43681-023-00314-4

Habib, S., Vogel, T., Anli, X., & Thorne, E. (2024). How does generative artificial intelligence impact student creativity? *Journal of Creativity, 34*, Article 100072. https://doi.org/10.1016/j.yjoc.2023.100072

Haraway, D. J. (2016). *A cyborg manifesto*. University of Minnesota Press.

Heidegger, M. (1977). *The question concerning technology and other essays* (W. Lovitt, Trans.). Harper & Row. (Original work published 1954)

Kim, N. J., & Kim, M. K. (2022). Teacher's perceptions of using an artificial intelligence-based educational tool for scientific writing. *Frontiers in Education, 7*, Article 755914. https://doi.org/10.3389/feduc.2022.755914

Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1–22). Association for Computing Machinery. https://doi.org/10.1145/3706598.3713778

Lin, H., & Chen, Q. (2024). Artificial intelligence (AI)-integrated educational applications and college students' creativity and academic emotions: Students and teachers' perceptions and attitudes. *BMC Psychology, 12*, Article 487. https://doi.org/10.1186/s40359-024-01979-0

Liu, G., Christian, B., Dumbalska, T., Bakker, M. A., & Dubey, R. (2026). *AI assistance reduces persistence and hurts independent performance* [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2604.04721

Livberber, T., & Ayvaz, S. (2023). The impact of artificial intelligence in academia: Views of Turkish academics on ChatGPT. *Heliyon, 9*(9), Article e19688. https://doi.org/10.1016/j.heliyon.2023.e19688

Malik, A., Khan, M. L., Hussain, K., Qadir, J., & Tarhini, A. (2025). AI in higher education: Unveiling academicians' perspectives on teaching, research, and ethics in the age of

ChatGPT. *Interactive Learning Environments, 33*(6), 2390–2406. https://doi.org/10.1080/10494820.2024.2409407

Marrone, R., Taddeo, V., & Hill, G. (2022). Creativity and artificial intelligence—A student perspective. *Journal of Intelligence, 10*(3), Article 65. https://doi.org/10.3390/jintelligence10030065

Mollick, E. (2024). *Co-intelligence: Living and working with AI*. WH Allen.

Nolen-Hoeksema, S., Wisco, B. E., & Lyubomirsky, S. (2008). Rethinking rumination. *Perspectives on Psychological Science, 3*(5), 400–424. https://doi.org/10.1111/j.1745-6924.2008.00088.x

Panke, S. (2025). How can (A)I research this? An autoethnographic exploration of generative AI in research, teaching and instructional design. *Journal of Teacher Education, 76*(3), 230–244. https://doi.org/10.1177/00224871251325065

Pensoneau-Conway, S. L. (2023). Autoethnography: Storied scholarship. In R. J. Tierney, F. Rizvi, & K. Ercikan (Eds.), *International encyclopedia of education* (4th ed., pp. 102–106). Elsevier. https://doi.org/10.1016/B978-0-12-818630-5.11012-7

Pisica, A. I., Edu, T., Zaharia, R. M., & Zaharia, R. (2023). Implementing artificial intelligence in higher education: Pros and cons from the perspectives of academics. *Societies, 13*(5), Article 118. https://doi.org/10.3390/soc13050118

Reed-Danahay, D. (1997). *Auto/ethnography: Rewriting the self and the social*. Routledge.

Reed-Danahay, D. (2009). Anthropologists, education, and autoethnography. *Reviews in Anthropology, 38*(1), 28–47. https://doi.org/10.1080/00938150802672931

Ringo, D. S. (2025). The effect of generative AI use on doctoral students' academic research progress: The moderating role of hedonic gratification. *Cogent Education, 12*(1), Article 2475268. https://doi.org/10.1080/2331186X.2025.2475268

Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. *Nature, 631*, 755–759. https://doi.org/10.1038/s41586-024-07566-y

Tian, J., & Zhang, R. (2025). Learners' AI dependence and critical thinking: The psychological mechanism of fatigue and the social buffering role of AI literacy. *Acta Psychologica, 260*, Article 105725. https://doi.org/10.1016/j.actpsy.2025.105725

Van Wyk, M. M., Adarkwah, M. A., & Amponsah, S. (2023). Why all the hype about ChatGPT? Academics' views of a chat-based conversational learning strategy at an open distance e-learning institution. *Open Praxis, 15*(3), 214–225. https://doi.org/10.55982/openpraxis.15.3.563

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems 30* (pp. 5998–6008). Curran Associates. https://doi.org/10.48550/arXiv.1706.03762

Vinchon, F., Lubart, T., Bartolotta, S., Gironnay, V., Botella, M., Bourgeois-Bougrine, S., Burkhardt, J., Bonnardel, N., Corazza, G. E., Glăveanu, V., Hanchett Hanson, M., Ivcevic, Z., Karwowski, M., Kaufman, J. C., Okada, T., Reiter-Palmon, R., & Gaggioli, A. (2023). Artificial intelligence & creativity: A manifesto for collaboration. *Journal of Creative Behavior, 57*(4), 472–484. https://doi.org/10.1002/jocb.597

Voita, E., Talbot, D., Moiseev, F., Sennrich, R., & Titov, I. (2019). Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 5797–5808). Association for Computational Linguistics. https://doi.org/10.18653/v1/P19-1580

Wallas, G. (1926). *The art of thought*. Jonathan Cape.

Wang, K., Li, J., Yang, S., Zhang, Z., & Wang, D. (2025). *When truth is overridden: Uncovering the internal origins of sycophancy in large language models* [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2508.02087

Wenger, E. (2024). AI produces gibberish when trained on too much AI-generated data. *Nature, 631*, 742–743. https://doi.org/10.1038/d41586-024-02355-z

Zhang, H., Xiong, X., Guo, L., Wang, X., Ye, T., Wang, X., & Ma, S. (2025). Posthumanist challenges and opportunities for teachers in the era of artificial intelligence. *Journal of Posthumanism, 5*(3), 673–683. https://doi.org/10.63332/joph.v5i3.775