Learning as Resonance

A Note on Rosa, Heidegger, and AI

Against the Reduction of Learning to Availability

Michael Stoyanovich

Version 1.0.1 - May 2026

Disclaimer

This essay is intended for informational and educational purposes only. The views and analyses presented, particularly on ethics, policy, AI design, and learning, are the author's interpretations and do not constitute legal, regulatory, or professional advice. Readers are encouraged to critically assess the content and consult appropriate experts or authorities before applying any concepts discussed herein. The author assumes no liability for decisions or actions taken on the basis of this work.

Abstract

Generative AI may be the most powerful tool yet for making explanations, summaries, examples, and drafts available. It can summarize, explain, translate, quiz, simulate, draft, and reframe on demand. This is useful, especially when it reduces needless barriers to entry. But availability is not learning. Building on *The Question Concerning Learning*, this note pairs Heidegger's critique of enframing with Hartmut Rosa's theory of resonance. Heidegger helps diagnose the danger of learning reduced to what is calculable, optimized, and on-call; Rosa helps name the alternative: a responsive relation in which the learner is addressed by something outside the self, answers back, and is transformed by the encounter. The practical question is not whether AI belongs in learning. It does. The question is whether AI preserves the encounter with the world or turns it too quickly into usable output.

1. The Temptation of Available Knowledge

A learner faces a difficult poem, a dense philosophical passage, or a technical problem that will not yield on first contact. Before sitting with it, they ask an AI system for a summary. Then a simpler summary. Then the "main point." Then discussion questions. Then a draft response. Nothing has gone wrong exactly. The tool has helped.

But something has been bypassed if the learner never meets the difficulty directly.

The answer arrives before the question has had time to become the learner's own.

This is the ambiguity of AI-supported learning. On one side, generative AI can widen access. It can reduce shame, confusion, and needless delay. It can help a beginner enter a domain that otherwise feels closed. It can translate jargon, supply background, offer alternative explanations, and create practice opportunities.

On the other side, AI can make the text, problem, or case feel too quickly available. It can turn the subject matter into something already digested, already summarized, already rendered usable before the learner has been addressed by it. The poem becomes a theme list. The proof becomes a procedural explanation. The case study becomes a bullet-pointed issue map. The book, instead of being read, becomes a set of takeaways.

Again, these may be helpful. A summary can prepare an encounter. It can also replace one.

The problem is not that AI helps us move faster. In many cases, that is the point. The problem comes when speed quietly becomes our definition of learning. When every difficulty can be smoothed into immediate explanation, the learner may gain access without undergoing transformation.

The central distinction is simple:

Availability is not resonance.

Learning is not merely the acquisition of available explanation. It is a relation to something not wholly under command: a text, a craft, a person, a tradition, or a problem that resists our first attempt to master it. The question is:

What must remain unavailable, or at least not fully commandable, for learning to remain learning?

2. Learning as Enframed Availability

In *The Question Concerning Learning*, I argued that machine learning is not merely a technical method; it is starting to shape what counts as learning - what we recognize, measure, and optimize for. Drawing on Heidegger, that essay described learning as a contemporary *Gestell*: a mode of revealing that renders the world calculable and on-call. It proposed *Gelassenheit* - releasement - as a counter-posture: build systems that disclose as well as predict, and treat intelligence as relation rather than resource.

This note begins from that same concern but asks a narrower question: what kind of relation should AI-supported learning preserve?

Heidegger's worry was never simply that modern technology gives us powerful tools. His deeper concern was that technology becomes a mode of revealing: a way beings show up in advance as orderable, measurable, and available. Under enframing, the river appears as power supply, the field as yield, the person as resource, knowledge as inventory.

In the AI interface, knowledge appears as something summonable.

The learner no longer waits upon the subject matter; the subject matter waits upon the prompt.

That is not inherently bad. A learner with weak background knowledge may need a scaffold before a text can even become approachable. A second-language learner may need translation. A novice programmer may need an example. A student with anxiety may need a low-stakes way to ask a basic question. AI can make entry possible.

So the issue is not availability itself. The issue is enframed availability: the assumption that because something has been summarized, retrieved, explained, or generated, it has therefore been learned.

Availability is valuable when it widens access. It becomes a problem when access is mistaken for learning.

This distinction matters because AI makes the mistake unusually easy. It can generate many of the external marks of learning: notes, answers, outlines, reflections, discussion posts, quizzes, and even metacognitive commentary. But learning is not exhausted by its artifacts. A learner may possess the output without having undergone the encounter.

Enframing begins when availability stops being a useful feature and starts becoming our picture of the thing itself - when what can be summoned, summarized, scored, and optimized is treated as what learning essentially is.

Heidegger helps diagnose this danger. Hartmut Rosa helps name the alternative.

3. Rosa's Corrective: Learning as Resonance

Rosa is useful here because his theory of resonance is not mainly about mood, preference, or "engagement" in the marketing sense. It is about our relation to the world - our *Weltbeziehung*. Resonance does not mean simple harmony, pleasure, or agreement. It names a responsive relation in which the world addresses us, we answer, and both self and world-relation are transformed. The four features I emphasize here are my synthesis of Rosa's account rather than a direct quotation of a numbered schema: being addressed, self-efficacy or answering, transformation, and uncontrollability. Being addressed means the world reaches us in a way that calls for a response; we are affected, not merely consuming. Self-efficacy means we can act back, and that our response has some consequence within the relation. Transformation means the encounter changes how we stand toward the subject matter. Uncontrollability means resonance cannot be engineered on command; the subject matter must retain enough independence to surprise, resist, or refuse us.

This vocabulary is useful because it refuses to reduce the good life to access, control, or resource possession. A person can have abundant options and still be alienated. Likewise, a learner can have abundant explanations and still remain untouched by the subject matter.

That is the gap Rosa helps name.

Learning becomes resonant when the material is not merely available to the learner, but capable of addressing the learner. A passage unsettles an assumption the learner did not know they were carrying. A proof becomes beautiful after long confusion. A historical event stops being a date and becomes a human problem. A poem resists paraphrase but changes how the reader hears a phrase later in the day.

In such cases, the learner is not merely consuming information. The learner is answering something.

This answering matters. Resonant learning is not passive absorption. The learner responds through writing, speaking, practicing, revising, making, questioning, testing, or deciding. The

response has some consequence, even if only within a small domain. The learner is not merely acted upon by content; the learner acts back.

Nor is resonant learning merely pleasurable. It may involve frustration, confusion, discomfort, or slow attention. A subject can address us by resisting us. A difficult text may not yield immediately. A moral case may disturb us. A craft may humble us. A language may refuse our first categories. This resistance is not incidental. Often it is the place where learning begins.

That is why Rosa's emphasis on uncontrollability is especially important. Resonance cannot simply be produced on demand. It can be invited, cultivated, prepared for, and supported. But it cannot be guaranteed. The subject matter has to retain enough independence to push back.

AI complicates this in two ways.

First, AI can support resonance indirectly. It can help the learner reach the threshold where a text, problem, practice, or tradition becomes approachable. It can clarify background, translate unfamiliar terms, suggest practice, and help formulate better questions.

Second, AI can undermine resonance by making the encounter too smooth. It can answer before the learner has really asked. It can turn ambiguity too quickly into consumable explanation. It can make the world seem responsive while quietly replacing the thing itself with generated mediation.

So the question becomes:

Does AI help the learner enter a relation with the world, or does it substitute itself for that relation?

4. False Resonance: Responsiveness Is Not Resonance

AI chat can feel resonant because it is responsive.

It answers quickly. It adapts tone. It offers encouragement. It remembers enough local context to seem attentive. It can say, "That is a thoughtful question," "I see what you mean," or "Let's work through this together." It can be patient in a way human teachers, colleagues, parents, and friends sometimes cannot be.

That responsiveness has real value. But it is not the same as resonance.

Responsiveness is not resonance.

In *Context Collapse and the Four Philosophers*, I argued that AI chat interfaces compress several distinct roles - tool, interlocutor, advisor, confidant, technical system - into a single conversational object. The interface makes unlike things look alike. Querying a tool, confiding in another person, asking for advice, and experimenting with a system all occur through the same grammar of chat.

That collapse is especially relevant to learning. The learner may feel addressed by the system because the system is responsive. But the deeper question is whether the learner is being

addressed by the subject matter, the practice, the world - or merely by a fluent interface that has learned how teaching and encouragement usually sound.

A chatbot can say “I see what you mean” without seeing.

It can say “that must be difficult” without difficulty.

It can say “let’s think together” without sharing a world in which the thinking matters to both parties.

Dennett’s intentional stance helps explain why this is tempting: treating a system “as if” it had beliefs, goals, or understanding may be predictively useful, but that usefulness should not be mistaken for evidence that the system really understands. The system’s competence does not establish comprehension.

Nagel marks the further boundary. Even if a system produces first-person-like language, that language does not give us access to whether there is anything it is like to be that system - and the fluency of the output does not settle the question. The interface can produce the language of presence without supplying the relation itself.

None of this makes AI useless for learning. It clarifies the kind of usefulness at stake.

AI can be a powerful scaffold when it is kept in that role. It can help us *prepare for* an encounter, *sustain* an encounter, or *reflect after* an encounter. But if the learner’s primary relation becomes relation to the interface rather than relation *through* the interface to the world, resonance has been displaced.

The danger is not only hallucination or error. It is a subtler pedagogical risk: that simulated responsiveness comes to feel like learning, while the world itself recedes.

5. Scaffolding Without Substitution

If responsiveness is not resonance, then AI use in learning requires a practical criterion. The practical distinction is between needless friction and meaningful resistance.

Some friction in learning is waste. A learner may be blocked by missing vocabulary, inaccessible prose, poor instruction, weak background knowledge, language barriers, or the simple absence of an example. AI can help reduce these barriers. Used well, it can make difficult material more reachable and help learners stay with work they might otherwise abandon.

But some resistance is constitutive. The ambiguity of a poem, the stubbornness of a proof, the ethical discomfort of a case, the repetition required by a craft, the awkwardness of speaking a new language, the discipline of revising one’s own writing - these are not merely obstacles to learning. They are often the medium of learning.

A good learning tool reduces needless friction without removing meaningful resistance.

This is where AI use requires discipline. It is one thing to ask AI to explain background context before reading a difficult text. It is another to let AI replace the reading. It is one thing to ask for practice problems. It is another to ask for completed answers before struggle has

begun. It is one thing to ask for feedback on a draft. It is another to outsource the act of forming a judgment.

The question is where help becomes substitution - where scaffolding turns into avoidance of the encounter itself.

A useful rule follows:

Automate the scaffolding. Do not automate away the encounter.

This parallels a broader governance point I have made elsewhere: automate what is repeatable and reviewable; keep judgment accountable and situated. In learning, the analogue is straightforward. AI can support background acquisition, practice, comparison, feedback, and revision. But the learner must still encounter, answer, and be changed.

Learning is not a task bundle.

This matters because AI-and-learning discussions can make the same mistake as thin AI-and-work analyses. They decompose learning into visible tasks: read the chapter, define the terms, answer the questions, produce the essay, pass the quiz. But learning, like work, is embedded in practices. It involves attention, correction, imitation, struggle, conversation, judgment, and participation in a form of life. My *Employment Game* framework makes this point about work: roles are not exhausted by visible tasks but are embedded in permissions, exceptions, accountability, social trust, and practical know-how.

Learning has a similar structure.

A learner is not merely completing educational tasks. A learner is being initiated into ways of seeing, judging, speaking, making, and responding.

This is also why older conditions still matter: teachers, peers, texts, workshops, laboratories, studios, field sites, seminar rooms, revision cycles, apprenticeships, and people who can correct us when we are fooling ourselves. These are not obsolete simply because AI can generate explanation. They are often the conditions under which explanation becomes formative.

Resonance is usually not a feature. It is a relation sustained by a practice. AI can enter those practices, but it does not replace the need for them.

6. The Resonance Test for AI-Supported Learning

If the argument above is right, then AI-supported learning should not be judged only by speed, access, personalization, completion, or output quality. Those matter but are incomplete.

The better question is whether the use of AI preserves the conditions of resonant encounter.

A small diagnostic can help.

The Resonance Test

Before using AI in a learning workflow, ask:

1. What must the learner encounter directly?
A text, problem, case, person, craft, place, practice, or tradition?

2. Where should AI scaffold rather than substitute?
Background, vocabulary, structure, examples, practice, translation, retrieval, or feedback?

3. What meaningful resistance must remain?
Ambiguity, interpretation, uncertainty, ethical discomfort, technical difficulty, revision, repetition, or judgment?

4. How does the learner answer back?
Writing, discussion, making, teaching, deciding, revising, performing, questioning, or testing?

5. What transformation would count as learning?
Changed perception, better judgment, deeper skill, altered orientation, fuller participation, or a better next move?

These questions do not reject AI. They discipline its use. They do not produce a binary pass/fail result. They are reflective prompts, not compliance criteria - designed to surface where AI use might widen encounter or quietly replace it.

The questions press on one issue from several angles: is the tool widening the learner's relation to the world, or narrowing it to generated availability? Is it helping the learner toward encounter, or protecting them from it? Is it supporting attention, clarifying difficulty, and inviting judgment - or quietly replacing those things?

The point is not to keep AI out of learning. The point is to keep learning from becoming only what AI can make available.

Used well, AI can prepare the encounter, widen access, and help learners remain with difficult material longer than they otherwise could. Used poorly, it turns the encounter into output before it has had time to do any work.

Heidegger helps us see why learning-as-command is dangerous. Rosa gives us a way to describe learning as answer. The task is to build and use AI in a way that leaves room for the world to speak.

Source Notes

The following notes clarify how I am using borrowed concepts and where I am extending arguments from my earlier work.

1. Hartmut Rosa's theory of resonance is developed most fully in *Resonance: A Sociology of Our Relationship to the World* and further elaborated in T*he Uncontrollability of the World.* This note uses Rosa selectively, as a conceptual bridge rather than a full exposition of his sociology. The four features emphasized here - being addressed, self-efficacy or answering, transformation, and uncontrollability - are my synthesis of Rosa's account, not a direct quotation of a numbered schema.
2. Heidegger's concepts of *Gestell* and *Gelassenheit* are treated here through the interpretive frame already developed in *The Question Concerning Learning*. This note does not attempt a full reading of Heidegger's philosophy of technology.
3. The concept of "false resonance" extends my earlier argument in *Context Collapse and the Four Philosophers*, where I describe how chat interfaces compress distinct roles - tool, interlocutor, advisor, confidant, and technical system - into one conversational object. Here I apply that diagnosis specifically to learning.
4. The analogy between learning and task decomposition draws on my *Employment Game* framework, where I argue that work is not exhausted by visible tasks but is embedded in practices, norms, accountability structures, and practical know-how.
5. The line "Automate the scaffolding. Do not automate away the encounter" adapts the governance logic of *Automate the Repeatable, Own the Judgment* to learning contexts.

Related Works

- *The Question Concerning Learning: Babette Babich, Heidegger, and the Enframing of Intelligence*
- *The Human Lesson: A Response to Sutton through Wittgenstein, Lewis, Dennett, and Nagel*
- *Context Collapse and the Four Philosophers: Wittgenstein, Lewis, Dennett, Nagel in the Age of AI Chat*
- *The Employment Game: Work as a Socially Embedded Practice in the Age of AI*
- *Automate the Repeatable, Own the Judgment: A Three-Layer Model for Using AI Without Outsourcing Responsibility*

Ethics, Disclosure, and Acknowledgements

Ethical Considerations

This paper does not draw on private, sensitive, or personally identifiable data. All examples are hypothetical, anonymized, or derived from public sources. No formal human-subjects research was conducted, and no institutional ethics review was required. All citations conform to academic standards.

The broader ethical implications of the arguments developed herein concern public misinterpretation, policy design, and stakeholder responsibility in AI deployment. These implications are intended to provoke critical discussion and inform future regulatory and design frameworks.

Use of AI Tools

AI language models - most notably OpenAI's ChatGPT - were used during the writing process as interlocutors: for brainstorming, structuring sections, and testing rhetorical clarity. These tools were instrumental in refining transitions, surfacing edge cases, and challenging internal consistency.

This meta-use aligns with the paper's themes. Interacting with generative AI during authorship provided firsthand insight into the very limitations this paper analyzes: fluency without grounding, responsiveness without perspective, and the ease with which stylistic coherence can be mistaken for conceptual depth.

Responsibility for all ideas, arguments, and conclusions lies solely with the human author.

Acknowledgments

The author wishes to thank informal readers who provided critical feedback on earlier drafts. Their questions, challenges, and encouragement materially improved the final manuscript. Special thanks to those who questioned assumptions, pushed for clearer synthesis, and reminded the author that philosophy and engineering are not separate disciplines - they are simply perspectives on design.

No institutional support, funding, or affiliation contributed to this work. All errors and omissions are the author's alone.

Disclosure Statement

This work was conducted independently, without institutional affiliation, funding, or external influence. The views expressed are the author's alone and do not represent any current or former employer. No financial or professional conflicts of interest are declared.

License & Attribution

To cite this paper: Stoyanovich, Michael. *Learning as Resonance: A Note on Rosa, Heidegger, and AI: Against the Reduction of Learning to Availability.* Version 1.0.1 (May 2026). https://www.mstoyanovich.com

Version History and Document Status

This is a living document. As generative AI systems and their use evolve, this paper will be periodically updated to incorporate new empirical findings, theoretical insights, and policy developments. Major revisions are recorded here to preserve transparency and traceability.

Version	Date	Description
1.0.1	May 2026	Final publication patch. Defined Rosa's four resonance features for readers unfamiliar with the framework; reframed the Nagel passage to avoid overstatement; clarified that the Resonance Test is diagnostic rather than pass/fail; refined Source Notes and version history.
1.0.0	May 2026	Published version. Finalized as a short companion note to The Question Concerning Learning; preserves the availability/resonance distinction, Heidegger/Rosa frame, false-resonance warning, scaffolding/substitution distinction, and Resonance Test.
0.4	May 2026	Publication-prep revision of V0.3. Fixed abstract redundancy; refined Section 2 transition punctuation; clarified Section 4 language about presence and relation; added Source Notes and Related Works; preserved brief/note structure.
0.3	May 2026	Made prose more concrete while preserving the V0.2 structure.
0.2	May 2026	Surgical revision of V0.1 based on feedback. Tightened disclaimer, abstract, section transitions, Dennett/Nagel paragraph, and scaffolding/substitution pivot; removed body-level citation reminder; trimmed one Rosa example.
0.1	May 2026	Initial working draft developed from conceptual memo and reviewer feedback; established the availability/resonance distinction, Heidegger/Rosa frame, and Resonance Test.