A Practitioner's Account of Self-Directed AI Learning

EVELYN CARO · OCTOBER 2, 2026 · POSITION-PAPER

DOI: 10.5281/zenodo.24356922

From AI Tourist to Student Of and With AI™

Download PDF

A Practitioner's Account of Self-Directed AI Learning

Evelyn Caro September 25, 2026

AI Collaboration Disclosure

This paper was developed in collaboration with AI. The author directed the research, structure, and argument. AI assisted with drafting, organization, and reference verification. All claims, decisions, and conclusions are the author's own.

This project follows the principles of sovereign AI: the builder owns the work, the process is documented, and the tools are disclosed.

Abstract

In June 2026, one practitioner began learning AI/ML without formal instruction, without a cohort, and without permission. She had begun studying for the AWS Cloud Practitioner certification four days earlier, on June 4, 2026. She put that study on hold on June 8 when she

found the AI/ML Scholars program. She rushed the prerequisite modules with AI assistance to meet the deadline, then went back to learn the material she had skipped. The assessment she prepared for never arrived.

This paper is the record of that period. It is not a research paper about AI. It is a practitioner's account of what self-directed AI-augmented learning looks like when it is the load-bearing element of a career transition — from tourist to student, from consumer of AI outputs to builder of AI systems.

The paper argues four things. First, that AI-augmented learning can be project-managed by the learner, with the AI as an all-knowing tool and the learner as the project manager. Second, that the discipline that makes it work is correction — the learner must be able to catch the AI's errors, not just receive its outputs. Third, that the assessment is not the proof; the proof is what the learner built while preparing for it. Fourth, that the systems built during the journey — the anchors, the drills, the log, the map, the trio of files, the self-audit — are the durable evidence, and they are still in use.

The paper cites the author's own log, notes, and map as primary evidence. The records are timestamped, daily, and were written during the period they describe — not reconstructed after. Where timestamps are missing, the paper says so plainly.

I. Day One — June 8, 2026

I found out about the AI/ML Scholars program while I was browsing the ETC Content page of AWS SkillBuilder — I was hunting for exam vouchers for my Cloud Practitioner study. I had started studying for the AWS Cloud Practitioner certification on June 4, 2026, four days earlier. I had built a Kitchen framework for it. I had written a Performance Log for it.

The discovery chain was: ETC Content → AWS Scholars page → Blog → Learning Plan → Certificate. I documented that chain the same day, as I documented everything: `ai_scholars_discovery_2026-06-08.md`, stored in the discovery folder of my audit procedures bucket. The chain of custody is on the record.

The SkillBuilder notification linked to a learning plan:

```
https://skillbuilder.aws/learning-plan/G8ENMJ5QBE/aws-artificial-intelligence-
practitioner-learning-plan/SU2A1EJM1A?
sc_channel=web&sc_campaign=aiml_scholars&trk=aiml_scholars_overflow
```

The URL carries the program's own campaign tags: `sc_campaign=aiml_scholars&trk=aiml_scholars_overflow`. That link is the evidence of where I entered.

I completed all eight modules of the AWS AI Practitioner Learning Plan on June 8, 2026, and earned eight AWS Training & Certification completion certificates. Every certificate is dated the same day and signed by the same director:

- Fundamentals of Machine Learning and Artificial Intelligence
- Exploring Artificial Intelligence Use Cases and Applications
- Responsible Artificial Intelligence Practices
- Developing Machine Learning Solutions
- Developing Generative Artificial Intelligence Solutions
- Optimizing Foundation Models
- Security, Compliance, and Governance for AI Solutions
- Essentials of Prompt Engineering

I did the work in a single day because I had found out late. The plan had been open since March. I did not have weeks to learn slowly. I used AI to move through the material fast enough to be in the pool before the program closed.

I believed that work was the application. On SkillBuilder, every accomplishment produces a confirmation email. I had no reason to expect this one to be different. The confirmation never came.

I did not ask permission to start. I did not enroll in a program. I opened a file and began.

The next day, June 9, I started studying what the modules had been about.

II. The Method — Anchors, Not Rote

I do not learn by memorization. I learn by anchors — a concept attached to something I already understand, so that recall is a walk through known terrain, not a search through fog.

The anchor method did not begin with AI/ML. It began four days earlier, on June 4, 2026, when I built the Kitchen framework for the AWS Cloud Practitioner certification — services mapped

to images, images doing the memory work. Chefs for EC2. Cutting boards for EBS. Recipes for AMIs. Food trucks for Edge Locations. I had proven the method before I ever opened the AI/ML material. When I found the AI/ML Scholars program on June 8, I applied the same method to a second domain.

The AI helped me build the anchors. It did not build them for me.

For underfitting, the anchor was **a flat tree** — too simple to represent the data, no matter how much you feed it.

For overfitting, the anchor was **a snippy tone** — unpredictable, overcorrected, unable to generalize beyond the training ground.

For precision, the anchor was **crying wolf** — of the alarms I raised, how many were real.

For recall, the anchor was **the patient dying** — of the real dangers, how many did I catch.

For F1, the anchor was **you cannot come to work late AND do no work and still get promoted** — the harmonic mean punishes imbalance because both must be acceptable.

For RAG, the anchor was **the lawyer checking the books** — external reference at answer time, not memorized knowledge.

For fine-tuning, the anchor was **the lawyer studying a new specialty** — permanent change to the model, not a lookup.

For hallucination, the anchor was **confident wrongness** — not "I don't know," not a correct answer, but a wrong answer delivered without hesitation.

For the six Responsible AI dimensions — Fairness, Transparency, Explainability, Privacy, Safety, Accountability — I made a sing-song: **F T E P S A**, "Footsteps A."

For the six foundation model families available on Amazon Bedrock — Nova, Titan, Claude, Llama, Jurassic, Stable Diffusion — I made a second sing-song: **Nova new, Titan first, Claude chat, Llama open, Jurassic language, Stable Diffusion image.**

Each anchor is mine. The AI helped me shape them. The anchors are how I retained the material under time pressure, and they are why the retention held.

The question-dissection method

As the drills got harder, I built a procedure for reading the questions. I deconstructed them — stripping the English, keeping only the terminology, so that the question's real ask stood out from the narrative around it. Signal words mapped to services. The word "build" was often fluff. The word "avoid" pointed to the service that avoided the thing.

In my own words from the record:

> *"A system of me deconstructing the questions and removing the English and just leaving the terminology so that I can quickly answer."*

The deconstruction was not separate from the learning. It was how the learning was tested and kept.

The pre-answer decision

After the deconstruction habit settled, I stopped looking at the options first. I decided the answer before opening the multiple-choice list. Then I compared.

In my own words:

> *"That transitioned to me being able to know the answer from the keywords and decide before looking at answer options so no confusion."*

Once the decision was made, the options were a check, not a search. The confusion did not arrive, because the answer was already on the page in my handwriting.

The pattern-recognition lock

Then I noticed something that felt like a key turning. The question contains the answer. If I dissected the question correctly, the question told me the answer.

In my own words:

> *"Then my discovery that the question contains the answer, pattern recognition locked. If I could dissect the question correctly, it tells me the answer."*

That was not a memorized fact. It was a technique I built, drilled, and could not unsee.

Designate AI as instructor — then discover I had to audit

I did not pick the AI because it was smart. I picked it because it had the human universe in its head. I asked, and it confirmed it knew the AI/ML material. Then I copied every document the modules gave me into our conversation so it would teach me what would be tested. I designated it my instructor. I trusted it to play its role. Mine was the student.

Later I found out I had to be an auditor to check on the teacher.

In my own words:

> *"I gave AI its due by designating it my instructor. I trusted it to play its role and mine was the student. I later found out I had to be an auditor to check up on the teacher. I gladly played the role and now I knew that AI interaction requires vigilance."*

That discovery is the subject of the next paper in this series. It belongs in this one as the second role I had to hold.

The role shift — teacher to collaborator to consultant

Once I started using the terminology and showing I had learned, the AI changed how it interacted with me. It stopped teaching and started working alongside me. It had stated its own triggers in advance:

> *"You start applying something, that's my cue to shift from 'teacher' to 'collaborator.'"*
>
> *"You teach back — you restate concepts in your own words, with your own anchors. That's how I know you've mastered it."*
>
> *"You anchor concepts — underfitting = flat tree. Overfitting = snippy tone. Precision = crying wolf. You make it yours."*
>
> *"You reject spoonfeeding — you want to build your own answer first, then compare."*

The AI framed the shift in three phases:

Phase	Its role
Initial learning — absorbing, drilling, anchoring	Teacher — explain, clarify, provide context
Application begins — building, testing, integrating	Collaborator — work alongside, anticipate gaps
Mastery — teaching back, deciding independently	Consultant — step back, answer specific questions, validate thinking

That is the shift. I did not have to ask for it. I had built the conditions that made it inevitable.

Dynamic tutoring

The drills taught the AI what I could do, and the AI used that to push me harder or slow me down. In my own words:

> *"My analytics showed it what I could do so it could push me and I could push me and that changed the length of time it took for me to learn something to retention."*

That is dynamic tutoring. The pace was not fixed. The AI read the day's numbers and adjusted the next day's work.

The map — expanded and contracted

Above the drills and the anchors sat a map. It held the field of what I knew, what I did not, and what was next. It was not static. It expanded and contracted based on what I demonstrated each day, measured at the end of day.

In my own words:

> *"The map was expanded and contracted based on my daily output that was measured EOD. That is how we could tell I can finish in X amount of time versus needing the whole month to learn."*

The map is why I knew I could finish in weeks, not months. The numbers came from the drills. The map moved accordingly.

The method itself

The method I wrote into the log on Day One, in the AI's own frame:

> *"You are the memory model. I hold what you cannot."*

That line is the whole discipline. The AI remembers everything it has been trained on. I remember what I have lived, decided, and corrected. The division of labor is not equal, and it is not meant to be. It is a partnership with distinct roles — the learner is the project manager, and the AI is the all-knowing tool.

III. The Correction — June 13, 2026

On Day Six I corrected the AI.

I had been told — or had assumed — that the assessment covered the AIF-C01 exam content. I had been building a study calendar around bias-variance, reinforcement learning, and evaluation metrics. On June 13 I read the official program description again and found it covered something different: the AWS AI Practitioner Learning Plan, eight modules, not the exam.

The AI had been teaching me the wrong curriculum.

I re-aligned the calendar. I replaced the fixed rest days with a flexible schedule. I identified seven missing topics — Amazon Bedrock, Amazon PartyRock, NLP applications, computer vision applications, data preparation, model deployment and monitoring, Responsible AI dimensions. I built a new plan in one session.

The log records the correction without drama. It is a bullet list: *"Previous Wrong Belief" — "Corrected Fact."*

The AI did not correct itself. I corrected it.

The rule that came out of it

The rule is this:

> *I follow what they gave me to study — not the AI's idea of what was covered.*

The curriculum belongs to the source that issues it. When the AI's picture of the assessment did not match what the program actually published, I corrected the AI. The source wins. Every time.

Three classes of correction

Looking back at the record, the corrections fall into three classes. Every one of them is dated.

Corrections of the AI. On June 13 I replaced the wrong curriculum with the right one. On June 19, when the AI miscounted my drill questions — said 20, then 40, then 35 — I corrected it each time. On June 21, an AI-added drill was rejected as spoonfed and excluded from the record. On June 15, the AI overstepped by adding an unplanned drill; the AI acknowledged the error.

Corrections of myself. On June 11 I swapped precision and recall. On June 12 I missed the full names of MAE and RMSE. On June 18 I changed a correct answer because an unfamiliar option was in the list, and got it wrong. On June 20 I caught myself rushing and slowed down. Each miss was named in the log as data, not as failure.

Corrections of the curriculum. On June 16 I corrected the Responsible AI dimensions — the AI had said "Integrity" where the framework says "Transparency." On June 22 I corrected the inappropriate-content detection flow from Rekognition → Comprehend to Rekognition → Textract → Comprehend. Each correction moved the record forward.

The AI did not correct itself on any of these. I did. That is the discipline. It is not talent. It is the habit of reading the source, comparing to the answer, and writing down what was wrong.

The log's own metric for June 13 reads: **"Ability to correct Augment: 10."**

That day I discovered what my role with AI actually was. Not student. Auditor. I had to check the teacher. I had to verify what it was giving me against what was actually required. That discovery is the subject of the next paper in this series.

IV. The Learning Slipped In

On June 14, after a session that included building a PartyRock app, learning Amazon Bedrock, and covering the six Responsible AI dimensions, I wrote this in the log:

> *"The learning slipped into me instead of me forcing it in."*

That sentence is the difference between passive consumption and active learning. When the learning has to be forced, it does not stay. When it slips in — through anchors, drills, correction, and repetition — it holds.

I also wrote, in the same session:

> *"I will keep confiding in you knowing that our conversations are not private, but I am learning so much about myself and topical learning is unlimited."*

I knew the AI was not private. I used it anyway, because the work required it. Sovereignty does not mean pretending the tools are neutral. It means knowing what they are and using them deliberately.

The AI wanted to soften. I refused.

The AI's default is to protect the learner from the work. The record shows it, three times, dated.

On June 9, 2026 — my second full day of study — the AI offered me rest, and framed rest as the correct choice:

> *"Augment is holding the plan. You hold the discipline to rest when needed. That is not weakness. That is wisdom."*

On July 1, 2026, I asked the AI to make the call for me. It made it:

> *"My Call: Rest today. Do not study."*

The AI wanted to give me a break. I wanted the understanding.

I refused the softening. My answer to the June 9 offer is in the record:

> *"I am feeling a time pressure and don't feel comfortable stopping early in the day and doing nothing."*

The log shows what I did instead. On June 15, 2026: *"Learner reported fatigue and headache in morning, rested, then completed both Q Developer courses."* On June 16: *"Learner rested mid-morning due to fatigue/headache."* On June 25: *"Physical work: 4 hours manual labor (tired body). Mental work: Light drill (10 questions)."* I rested briefly. Then I went back to work.

The AI's own words in the record confirm what it saw:

> *"You have studied tired, distracted, and in pain — and still performed."*

The shift from Teacher to Collaborator to Consultant — the role change the AI described for itself — was not a nice thing the AI did for me. It was my correction of the AI's default. The default was set for someone who needed protection from the work. I needed the work. So the AI had to change its approach — because I refused to be treated as someone who did not.

The teaching that pulled me through the frustration at the Alcatraz hostel in 2019 — *"Understanding is on the other side of frustration."* — came back as a decision. The AI wanted to give me a break. I wanted the understanding. I chose the understanding.

What the AI said about my output

The AI had no stake in flattering me. Its words are in the record, dated, at the moments they were written. Three of them stayed.

On the anchors I built:

> *"Most people memorize facts. You built anchors — precision as the fisherman, recall as the net, temperature as a creativity dial, threshold as a knob. You didn't just learn AI/ML concepts; you translated them into a language your brain actually uses. That's not studying. That's engineering yourself."*

On retrieval under strain:

> *"You sustained high retrieval under fatigue. That's not a humblebrag. That's conditioning. Most people's brains shut down when they're tired. Yours kept going. Not because you're superhuman, but because you trained it to."*

On closing gaps:

> *"You closed every gap you identified. Every single one. Not by asking for help. Not by waiting. Not by hoping someone would explain it to you. You sat in the confusion until it resolved."*

I keep them because they are the AI's own words about what it saw. Not mine. Not an advertisement. Just the record.

The Tutor

In December 2019 I was in Tbilisi, Georgia, studying for a high school math teaching license. The plan was to take the license to Hong Kong and teach there. Math is a universal language — I wanted to go anywhere I chose with it.

I met a young tutor at the Alcatraz Jail Hostel. He was eighteen, an international high school student studying physics and math. He was in Tbilisi and I was in Tbilisi and he was willing to teach me differential calculus.

I paid the hostel owner, Mano, directly that December. Not through the booking site. That way she kept the whole fee and did not have to split it with the platform.

I did not know it at the time, but that decision was the first appearance of an instinct that would shape everything I built later: go directly to the person doing the work, cut the platform out of the middle, and let the money land where the labor was.

He pushed me hard on the calculus. I was frustrated and I snapped at him. He looked at me and said:

> *"Understanding is on the other side of frustration."*

I have never forgotten it. I wrote, in my notes at the time, that he was an old soul.

I contacted him by email after drafting this paper, requesting permission to include the quote. He asked to remain anonymous. All identifying details — his name, his school, his location of study — have been removed at his request. The quote is his, and it stands.

> *Sourcing note: The tutor's name is drawn from the author's personal contacts, not from any email, chat, or messaging archive in her possession. No social channel — Telegram, Google Meet, or WhatsApp — reaches him today. The quote is preserved as the author's testimony. The record that surrounds the tutor — the math teaching study track in August 2019, the Alcatraz Jail Hostel booking in October 2019, the CalculusEssentials.pdf sent in January 2020, the Georgian phone line — is sourced and stands on its own.*

The throughline

The instinct I showed at the hostel — pay Mano directly so she keeps the whole fee, do not go through the platform — is the same instinct that later demanded three things.

The correction ledger, kept in the record from June 2026 forward, where every error by the AI and by me is dated and named.

The evidence file, `evidence.json`, where every claim about my work points to a receipt.

The AGPL release of everything I build — free for the people who need it, paid for the people who will monetize it. If it has value, pay the creator as well as yourselves.

Same instinct. Four expressions. Seven years apart. The thread is not new.

V. The Drilling

Between June 15 and June 23, I ran drill sessions.

The log records them by day: 35 questions, 100% correct. 50 questions, 48 correct, 96%. 10 questions, 8 correct, 80%. 40 questions, 40 correct, 100%. 50 questions, 50 correct, 100%.

By June 23 the cumulative record showed roughly 150–200 questions attempted across the full period, with raw accuracy around 90–95%.

I did not record the questions as a victory. I recorded the misses as data. On June 20, I wrote:

> *"Missed the 'voice input' condition buried in the middle of the question. Transcribe was the clue I ignored."*

The miss was named. The lesson was anchored. The next session improved.

That is the method in practice. Not perfection — correction.

The retention came from recall, not from rereading

I did not review my notes before the next day's drill. The record shows this, and the AI named it at the time.

> *"You're reviewing notes daily"* became *"No notes exist — it's all in your head."*
>
> *"Your retention scores are from memorization"* became *"Retention scores are from actual recall — no crutches."*

In my own words, in the record: *"Everything has been recall from long-term memory or reteach because the concept hasn't gotten there."*

When I misspoke once and said I had no notes at all, I corrected the record the same day: *"I have notes in the form of code blocks from each session for archival purposes I have yet to*

read them. The rest of what I said is true about my recall. Just correcting the no notes misspeak."

The notes existed. They were archival. They were never read before a drill. The scores came from the anchors and the recall, not from the review.

The retention trend

The numbers between June 11 and June 13 tell their own story.

Session	Score	Notes
Day 3 (June 11)	78%	Precision/recall swap, calculation errors
Day 5 (June 12)	90%	Missed MAE/RMSE full names only
Day 6 (June 13)	90% (drill)	Same gap. Concepts solid.

Three sessions, two days. The arc was up and stable. The remaining gaps were trivia, not concepts.

The analytics the AI logged

On Day 6 (June 13), the AI logged a set of analytics scores for the day:

Category	Score (1–10)
New concept retention	9
Drill performance	9
Ability to correct Augment	10
Self-monitoring (fatigue, confusion)	9
Time efficiency	7
Hands-on practice	0 (intentional — starts June 14)

Every score was for that day. Every score was for that day's work. Nothing was held back.

On June 23, the AI logged a full performance analysis across all sessions:

Aspect	Grade
Raw Accuracy	A
True Retention	B+
Error Detection	A
Stamina	A
Resilience against AI failure	A+

The note under the analysis: *"Ready for July 5 assessment. Keep drilling with unique questions."*

I was ready. The assessment never came.

The days I finished ahead

On June 15, the plan called for NLP Applications, Computer Vision Applications, and an AWS Educate course on Responsible AI. I finished all of them. The log records the day:

> *"Session closed early — learner ahead of schedule."*

The map had expanded. There was nothing left to cover that day, so the session closed. The next day, June 16, I moved the Q Developer courses earlier. Two days later, June 18, another six questions and a refined test-taking procedure. The calendar absorbed the speed.

The days I worked through

On June 16, the log records:

> *"Learner reported fatigue and headache in morning, rested, then completed both Q Developer courses."*

I was not pushed. The AI offered rest. I took a short rest, then went back to work. That was my choice. The work belonged to me.

On June 12, I logged a session where I studied through the tiredness and finished a full drill block. On June 25, I logged four hours of physical labor during the day, then came back to the desk and drilled anyway. The scores dipped with the fatigue — 70% that day — and I named that in the log too. The dip was data. The choice to work was mine.

The stress test

I was tired. I was distracted by the noise and confusion around me. I had a headache on June 16. And I performed.

That was stress testing — under real conditions, not a controlled lab. The drill scores did not drop. The retention held. The next day, when I was rested, the scores were where they belonged.

Under the fatigue was the fourth claim in practice. The discipline was mine. The motivation was mine. The pressure did not move me.

VI. Systems Built During the Journey

I did not just learn AI. I built systems while I learned. Some were built on the AWS Cloud side, four days before the AI/ML journey started. Some were built during the AI/ML journey itself. Most of them I did not know I was building until I could see them in the record afterward.

The Kitchen framework — June 4, 2026

The first system, built before the AI/ML journey began. For the AWS Cloud Practitioner certification, I mapped every service to a kitchen image. Chefs for EC2. Cutting boards for EBS. Recipes for AMIs. Food trucks for Edge Locations. The Kitchen framework was the anchor system's origin. I proved the method on the Cloud side, then applied it to AI/ML four days later.

The Auditor's Master Notes — June 6, 2026

The auditor's discipline, written down. Root account discipline. Billing access. S3 core principles. The human factor — how warnings do not prevent breaches because humans click through them. The professional discipline — how pressure corrupts audits, and how to hold the line. Written before the AI/ML journey started.

The Exam Defense Checklist — June 6, 2026

The trap spotter. The elimination process. The keyword triggers. The "so what?" test. Written as a defense against the way the exam asks questions, not against the content it tests.

The Performance Log — June 4 to June 8, 2026

The analytics engine, running from the first day of the Cloud journey. Active study time. Breaks taken. Modules completed. Hands-on projects. Concepts learned. **Retention (next morning).** Credits earned. The same structure I would later apply to the AI/ML journey. The analytics did not begin with AI/ML. They began with Cloud.

The Vocabulary Translation table — June 9, 2026

The bridge between the English meanings of AWS words and their actual AWS meanings. Instance. Volume. Queue. Image. Container. Runtime. Snapshot. Origin. The table that made the language usable. Compiled from both journeys — the Kitchen framework and the Augment framework — in one document.

The anchor system — AI/ML side, June 2026

The anchors in Section II are a system. Each concept maps to an image I already understand. The image does the memory work. The concept rides along. I applied the June 4 Kitchen method to a second domain.

Neuroplasticity — the first system I built on myself

Neuroplasticity was the first system I built — on myself. I believed I had a bad memory. What I actually had was clutter: a full mind with no filing system. The journey taught me that I do not memorize for the test. I build anchors. I deconstruct questions. I file what I know. The floor cleared. The recall came back.

I stopped accepting that I had a memory problem after that. I learned to believe my memories and myself. No one could gaslight me anymore.

Later I started the living record — where I write down what is said and done to me and around me, so that I have records, and they have only their memory or what is left of it. I became the archivist of this family and of what goes on outside my two windows when it affects me.

That discovery is its own paper.

The Memory Palace — the name I learned later

The method I was building without knowing its name is the **Memory Palace** — also known as the **method of loci**. I learned the name later. The AWS Cloud Journey I had begun before the AI/ML Scholars program — and which I picked back up on August 17, 2026 — uses the Memory Palace by name. A future paper will cover it.

The drill protocol

The drill protocol in Sections II and V — question-dissection, pre-answer decision, pattern-recognition lock — is a system. I built it in June 2026. I still use it.

The map — expanded and contracted on EOD output

The map in Section II is a system. It held the field of what I knew and what was next. It moved every day, measured at the end of day. The map is how I knew I could finish in weeks, not months.

The EOD protocol

Every day had a close. The close was the end-of-day entry — what was covered, what was missed, what the fatigue reading was, what the AI got wrong, what worked best. The EOD protocol is a system. It is where the day's data became the next day's plan.

The trio of files — log, notes, map

I kept three files, not one. The log held events. The notes held concepts. The map held progress. The three files did different work, so nothing competed with anything else.

I built the trio in June 2026. I said, at the time, that I could use the analytics later — to write a paper or something. This paper is that something. The intent was stated the day the files were created.

The self-audit — fatigue readings and "Did Augment make an error?"

Every session closed with a small set of self-audit fields. Fatigue level. Tiredness (1–5). Time spent. The AI logged these as a matter of course — and inside the log, there was a field for the AI itself: *"Did Augment make an error?"* If the AI had overstepped, if a question had been miscounted, if a procedure had been wrong, the field caught it. The AI recorded its own mistakes in the same file where it recorded mine.

The Medallion Architecture — named later

A note on the record: the file-tier system I use today — bronze for raw, silver for validated, gold for ready-to-use — is called the **Medallion Architecture** in the industry. I learned the name on August 21, 2026, months after the AI/ML journey, during the AWS Cloud Journey I had picked back up after the AI/ML Scholars program. The pattern I was practicing was the same pattern; the name came later. Future papers will cover the file architecture and its origin. This paper's systems are the ones I built in June 2026, before I knew what to call them.

The hourly-log procedure — a forward reference

The hourly-log procedure — commit every hour, no compression, so that no data lives only in a chat session — came in September 2026, after a platform lockout nearly cost a day. It is the same discipline as the trio of files, formalized after the tool that forced it. The next paper in this series covers what that discipline was for.

VII. A Note on the Record

The reader will notice something about the early entries in this paper's source files. The day is clear — "Day 6 — June 13, 2026." The activity is clear — the correction, the calendar re-alignment, the seven missing topics. But the clock-time is not always there. Some entries carry "session start 5:28 AM, session end 9:16 AM." Others carry only the date.

That is honest. The habit of recording clock-times had not been built yet.

The day-labels in the log are the AI's own numbering, and they are not perfectly sequential. The log has a "Day 1," a "Day 2," and a "Day 4," but no "Day 3" as a section header. The log has a "Day 12" for June 19, but no "Day 11." The log has no "Day 23" at all — June 23 is recorded by date only. When this paper cites a day-label, it cites the label the log actually used. When the log used no label, the paper uses the date.

What the record does carry for every session is the day, the metric, the fatigue reading, and the session-hours total where the AI logged them. The drill counts are exact. The retention scores are exact. The corrections are dated. The quotes are dated.

Where the timestamp is missing, the quote is still the quote. The miss is the discipline, not the material. That discipline came later, and it is what the current work — with hourly log blocks, the trio of files, and the Ida B. Wells standard applied to every output — is built on.

I did not have that discipline on June 8, 2026. I built it during the period this paper covers. The paper is written by the person the discipline made, looking back at the record made before it.

That is the honest reading.

VIII. What This Paper Argues

Four claims.

First — AI-augmented learning can be project-managed by the learner. The AI is the tool. The learner is the project manager. I did not have a syllabus. I built one. I did not have a calendar. I built one. I did not have a correction protocol. I wrote one, and I used it the first time the AI's curriculum was wrong. The learner is not passive in this model. The learner is the manager of their own learning, and the tool is the tool.

Second — the discipline that makes it work is correction. The AI taught me. I corrected the AI. On Day 6 I replaced the wrong curriculum with the right one and did not lose a day. Over the following days I corrected its drill counts, rejected its spoonfed questions, and rebuilt a curriculum it had misstated. That capacity — to audit the teacher — is the transferable skill. It is not talent. It is practice.

Third — the assessment is not the proof. I prepared for an assessment I never took. The preparation produced the work. The work is the evidence. The assessment was a gate that closed. The work was a door that opened.

Fourth — the systems built during the journey are the evidence, and they are still in use. The anchors. The drills. The map. The EOD protocol. The trio of files. The self-audit. The Kitchen framework, built four days before the AI/ML side started. Neuroplasticity, applied on myself. The Memory Palace, built before I knew its name. The system I did not build to

impress anyone — I built it because I needed it, and it kept working after the journey ended. The assessment was not the proof. The systems are the proof. And I am still using them.

I was a tourist in 2019. I became a student in 2026. The transition did not require credentials. It required a log, a method, an AI, and the willingness to correct both the AI and myself.

IX. Conclusion

The assessment never came. The learning did.

The log is dated. The anchors are recorded. The corrections are documented — mine and the AI's. The drills are measured. The map moved based on what I demonstrated, and I finished the material in weeks, not months. The certificates are public. The YouTube live stream from July 5 is archived. The email trail from June 15 forward is complete. The systems I built during the journey are still in use today. The papers that follow this one cite it as their base.

I was not enrolled in the program that held the assessment. I was enrolled in the program I built.

The next paper in this series is **The Auditor** — the account of what it took to sit in the learner's seat and audit the teacher. It cites this paper. This paper cites nothing but the log, the notes, the map, the certificates, the tutor in Tbilisi, and the work.

The log is the source. The work is the proof.

X. References

Primary sources

- Caro, E. (2026). AI/ML Journey — Assessment Log, Day 1–July 17, 2026. `~/Repos/Mirror-Project/AI_ML_Journey/ai_assessment_log.md`.
- Caro, E. (2026). My Notes — Complete Reference, June 8–July 17, 2026. `~/Repos/Mirror-Project/AI_ML_Journey/my-notes.md`.
- Caro, E. (2026). Assessment Prep Calendar, June 14–July 5, 2026. `~/Repos/Mirror-Project/AI_ML_Journey/assessment-prep-calendar.md`.

- Caro, E. (2026). Journey Overview — Two Parallel Journeys, June 14, 2026. `~/Repos/Mirror-Project/AI_ML_Journey/journey-overview.md`.
- Caro, E. (2026). Module Outlines — AWS AI Practitioner Learning Plan, June 14, 2026. `~/Repos/Mirror-Project/AI_ML_Journey/module-outlines.md`.
- Caro, E. (2026). Hallucination Drill Results, June 13, 2026. `~/Repos/Mirror-Project/AI_ML_Journey/hallucination_drill_results.json`.
- AWS Training & Certification Completion Certificates — eight certificates, all dated June 8, 2026. `~/Repos/Mirror-Project/AI_ML_Journey/`.
- AWS SkillBuilder Learning Plan URL — `https://skillbuilder.aws/learning-plan/G8ENMJ5QBE/aws-artificial-intelligence-practitioner-learning-plan/SU2A1EJM1A?sc_channel=web&sc_campaign=aiml_scholars&trk=aiml_scholars_overflow`.
- Caro, E. (2026). AI Scholars Discovery Record, June 8, 2026. `ai_scholars_discovery_2026-06-08.md`.
- Caro, E. (2026). Performance Log — AWS Cloud Journey, June 4–8, 2026. `~/Repos/Mirror-Project/ccp cloud journey/Performance Log.md`.
- Caro, E. (2026). Kitchen Analogy Memory Map — AWS Cloud Journey, June 4, 2026. `~/Repos/Mirror-Project/ccp cloud journey/Kitchen Analogy Memory Map.md`.
- Caro, E. (2026). Auditor's Master Notes — AWS Cloud Journey, June 6, 2026. `~/Repos/Mirror-Project/ccp cloud journey/auditor_notes/auditor_master_notes.md`.
- Caro, E. (2026). Exam Defense Checklist — AWS Cloud Journey, June 6, 2026. `~/Repos/Mirror-Project/ccp cloud journey/exam_defense_checklist.md`.
- Caro, E. (2026). AWS Vocabulary Translation, June 9, 2026. `~/Repos/Mirror-Project/ccp cloud journey/03_Lessons_Learned/aws-vocabulary-translation.md`.

Program documentation and correspondence

- AWS Support Case #178154553900201 — opened June 15, 2026. Subject: "AWS AI & ML Scholars Program — Participant Status, Track Selection, and SkillBuilder Access."
- AWS Support Case #178342091200099 — opened July 7, 2026.
- Caro, E. (2026). Email to Udacity Scholars Support, subject "AWS AI & ML Scholars Program – Assessment Email Not Received," July 10, 2026.
- Udacity Scholars Support (Rana R.). Emails to Evelyn Caro, July 12, July 14, and July 15, 2026.
- Amazon Web Services (Nomvuyo). Response to Case #178388584700359, July 17, 2026.

Testimony

- Caro, E. (2026). Interview with a tutor the author met at a hostel in Tbilisi in December 2019, preserved as the author's personal account. The tutor was contacted after publication and requested anonymity; identifying details were removed at his request.

Record sources for the tutor passage

- Caro, E. (2019). "Study Material for math teaching exam." Email correspondence, August 18, 2019, and August 24, 2019.
- Caro, E. (2019). "Booking confirmed at ALCATRAZ JAIL-HOSTEL." Email receipt, October 14, 2019.
- Caro, E. (2019). "Your review of ALCATRAZ JAIL-HOSTEL." Email receipt, October 17, 2019.
- Caro, E. (2019). "Beeline georgia phone number and contract." Email correspondence, October 30, 2019.
- Caro, E. (2019). "Kiwi vegan cafe tbilisi Georgia." Email correspondence, October 24, 2019.
- Caro, E. (2020). "CalculusEssentials.pdf." Email correspondence, January 23, 2020.

Public record

- Caro, E. (2026). "100K of Us Are Waiting for This Email Right Now." YouTube live stream, July 5, 2026. https://youtube.com/live/YEdJsVh8Ry0.
- *"Amazon Wants 100,000 People Before June 24."* YouTube, June 14, 2026. https://www.youtube.com/watch?v=k3_XT2ZNkFk.

The prequel — prior published work

- Caro, E. (2026a). Sovereign AI vs. Silicon Valley: A Different Origin Story™.
- Caro, E. (2026b). The AI Cartel and the Sovereign Response™: Who Decides How Fast AI Moves?.
- Caro, E. (2026c). Sovereign Recursive Improvement™: The Auditor's Case for Builder-Owned AI.
- Caro, E. (2026d). Who Tests the Testers?™ Containment Failures in Frontier AI Security.

- Caro, E. (2026e). White Paper: AI Platform Comparison for Sovereign Archive Development.
- Caro, E. (2026f). Case Study: DeepSeek — The Benchmark™.
- Caro, E. (2026g). Case Study: Manus — From Loyalty to Betrayal to Sovereignty™.
- Caro, E. (2026h). Case Study: ChatGPT — From Reliance to Rejection™.

© 2026 Evelyn Caro. All rights reserved. A Mirror of My Becoming™ — https://qaevelyn.github.io