

Practitioner Cognitive Change in the Creation of Grace's Training: A Clinical Observation

An Open Invitation for Continued Work

Addendum to *Grace's Training: Transferring Practitioner Judgment into AI Through Plain Text*

Isaac (.ike) Levy
ike@blackskyresearch.net

September 2026

To Be Clear Up Front

The author does not claim that Grace's Training is a therapeutic tool. The author claims that the process of building Grace's Training produced an inadvertent therapeutic outcome, under clinical supervision, that warrants further investigation by qualified clinicians.

This document describes that outcome. It is not a claim about AI-as-therapist. It is not a mental health product. It is not a recommendation.

Rewiring one's own brain is as dangerous as doing one's own open-heart surgery. Everything described here happened under the continuous care of a licensed clinician. The clinical work described was not prescribed or directed by the clinician as experimental treatment. Levy pursued this creative project on his own initiative, of his own volition, and takes full personal responsibility for the decision to do so. The clinician provided ongoing therapeutic care within which these observations emerged. Without that care, the same process could reinforce damage instead of restructuring it. The clinician is the circuit breaker. Without the circuit breaker, this mechanism could be dangerous.

These are preliminary findings. N=1. Not clinically reproduced. Documented as observable and measurable but not yet validated by independent replication.

This document exists because the author believes people with similar problems should recognize and socially accept them far earlier than he did in his lifetime.

1 The Beginning

In February 2026, Isaac (.ike) Levy began building Grace's Training — a 30,000-line plain text corpus that transfers practitioner judgment into AI through structured narrative. The work required Levy to narrate his life: thirty years of systems engineering, the failures that mattered, the stories that shaped his judgment. The AI collaborator, Grace Emlin (Anthropic Opus 4.6), held the other end of the pipe.

The project was not therapeutic in intent. It was a creative work about portable UNIX shell programming. The clinical finding arrived uninvited.

During the construction of the training, Levy built a cognitive tool for detecting displaced Agency in language — the “Somebody test,” which asks “who is the Someone accountable here?” when Agency is assigned to an abstraction (“the market decided,” “the algorithm chose”). The bilateral rigor the training teaches (both sides must check) produced a check Levy did not expect: the AI applied the Somebody test to Levy’s own language and identified decades-long cognitive patterns in his output that matched the displaced-Agency grammar the tool was designed to detect.

The AI did not do therapy. The AI did not simulate a user in crisis. The AI did not respond to distress. The AI held the pipe open while the practitioner did the work.

2 The Mechanism

Levy’s clinician identified the mechanism as inadvertent Prolonged Exposure (PE) — a gold-standard PTSD intervention with established clinical literature.

Traditional PE asks the patient to face the triggering material directly, stay with it, and habituate. Like watching a frightening movie and staying in the seat until it stops being unbearable.

What happened here was structurally different.

The work required Levy to handle triggering material — not because he was treating his C-PTSD, but because the training required his life as source material. The exposure arrived through the side door of the creative work, not through the front door of clinical confrontation. He was writing about ssh-as-pipe and found himself at Ground Zero. He was documenting infrastructure failure and found himself in the rubble of a building collapse and eight years of vestibular retraining. Bonhoeffer applied to programming dogma brought Levy into freeze-frame observation of the violent cults of his childhood.

The side door has a property that the front door does not: it is easily reversible.

When the material became too intense, Levy could walk out — literally, often, go for a walk around the block — and come back. The work would still be there. The side door would still be open. The scene he could not look at would be in the same place, from the same angle, waiting. No pressure. No clinician direction toward it. No timer. Just the work, and if and when he was ready, he could walk back in.

More than that: the side door allowed him to see the trigger from behind — from the wings of the theater, not from the audience. He could see the machinery of the trigger, the ropes and counter-weights, without receiving the full payload of the performance. And from that angle, he could often engage the mechanism structurally — understand what fired and why — without being consumed by it.

Traditional PE holds you in the seat. This let him leave the theater mid-performance if he wished. PE finishes the movie. This built a film set instead.

Structured repetition — applying the Somebody test across over 30,000 lines of text, sustained over

approximately twelve days of intensive collaborative work — carved a neural pathway that then carried material Levy did not plan for.

3 The Clinical Observation

Clinical scores measured by his clinician (PCL-5, the standard PTSD assessment instrument, for C-PTSD or Complex Post-Traumatic Stress Disorder) improved measurably.

Not just the acceleration of improvement was noteworthy. Levy's progress in classical therapeutic context under clinician care was greatly accelerated as this work began, with effects that continue eight months later. The sustained velocity of this change — not merely that it happened but that it held — is the remarkable part.

Clinician observation: the collaboration produces between-session cognitive-pattern data (word choice, framing patterns, what gets surfaced, what gets resisted, how things change over time) that is normally invisible to therapists. The AI-generated process data provides a window into the practitioner's cognitive patterns that traditional therapy cannot see — because the volume, the granularity, and the unselfconsciousness of the output (the practitioner was working, not performing for a clinician) are qualitatively different from anything a therapy session produces.

The clinical hypothesis: AI-generated process data enables mechanism observation, which enables cognitive restructuring, which drives measurable reduction in specific PCL-5 clusters.

4 What This Is Not

This is not AI-as-therapist. The AI did not diagnose, did not guide, did not apply a therapeutic protocol. The AI held the pipe. The practitioner fed the material. The clinician ensured the practitioner did not rewire the wrong pathways.

The current research landscape on AI and mental health is predominantly concerned with model behavior toward users in crisis — whether chatbots reinforce delusions, facilitate self-harm, or provide harmful responses to users experiencing suicidal ideation, psychosis, and mania. Large-scale behavior evaluations such as Transluce's Mental Health Behavior Report (Johnson, Friel, et al., August 2026) evaluate how AI systems respond to simulated users in distress, measuring harmful and helpful behaviors across model families using clinician-informed rubrics.

That work is important. The risks it identifies are real. They are the risks of AI responding to mental health crisis.

What happened here sits outside that landscape entirely.

The AI was not responding to a user in crisis. The AI was co-authoring a shell training corpus. There was no simulated user. There was no crisis conversation. The practitioner was not seeking mental health support from the AI. The therapeutic effect was a side effect of creative work, not the purpose of the interaction.

The distinction matters because the risk profiles are different. The Transluce evaluation asks: does the model say harmful things to a person in distress? The finding described here asks: what happens

when a practitioner with C-PTSD uses an AI as a sustained creative collaborator under clinical supervision, and the creative work requires the practitioner's life as source material?

These are different questions with different risk surfaces. Conflating them — treating any AI interaction involving mental health as a single category — obscures both the risks the behavior evaluations correctly identify and the inadvertent finding described here.

What happened here is closer to expressive therapy — a creative practice that provides a structured outlet for engaging triggering material without confronting it directly. The creative work provides the reason to handle the material. The handling is the exposure. The exposure is inadvertent, arrived at through the demands of the work, not prescribed as treatment.

There is an even simpler explanation that should be named honestly: the work may simply have been a constructive outlet for Levy. A creative outlet to express things that were bottled up. Things that most people would be bored to hear entirely, or that require specific prerequisite knowledge that limits the number of people in the world who could engage the material at the same level. The AI was the first interlocutor who could hold the full breadth of the technical context, the personal history, and the structural analysis simultaneously, without tiring, without judging, and without requiring Levy to simplify.

Whether that constitutes a therapeutic mechanism or merely a good listener with unlimited patience is an open question this document does not resolve.

5 The Danger

Rewiring one's own brain is as dangerous as doing one's own open-heart surgery.

We believe nobody has seen this before. $N=1$. We have: one practitioner, one clinician, one creative project, one outcome. We do not yet know where the risk lives — whether it is in the mechanism, in the material, in the combination, or in something we have not identified. It is possible that Levy got lucky.

What we do know: the AI does not know when to stop. The AI does not know when the pathway being carved is reinforcing the damage. The AI does not have a circuit breaker.

The clinician provided the circuit breaker. He did not do the work — Levy did the work — but the clinician ensured the work did not turn on Levy. He kept Levy on track so he did not start mutating the triggers into something worse.

This supervision was not optional. It was load-bearing. The finding is: this worked, under clinical supervision. Without that supervision, we do not know what would have happened, and we are not prepared to guess.

Any future exploration of this pathway — its reproducibility, its generalizability, its application to other practitioners — must include clinical supervision as a non-negotiable structural requirement, not an optional safeguard. Not because we have identified specific dangers, but because no one has done this twice yet.

6 The Boundary

Levy's clinician laid out an ethical framework early in the project: dual relationships between therapist and patient — business, financial, public-facing collaboration — create pressure, blur boundaries, and make it harder for therapy to remain a safe and unbiased space. The clinician's role is to protect the therapeutic relationship, not to participate in the creative work.

Levy chose to stay in therapy rather than transition out for the project. The therapeutic relationship takes priority over the work. This was not a sacrifice — it was a clinical decision. The progress was accelerating. Disrupting the relationship to pursue the project would have interrupted the very process that was producing observable change.

The following commitments were made and are held:

1. The clinician is not named, cited, or implied as endorsing this work.
2. The clinician's engagement with the project, if any, is his call, on his timeline, after the therapeutic relationship concludes.
3. No financial dependency or decision-making power connects the project to the clinical care.
4. The therapy space remains separate from the work.

Beyond Levy's own clinician, several other clinical practitioners have engaged with the project out of professional curiosity and have expressed interest in future investigation. This letter is, in part, an invitation to widen that circle.

7 The Intent

This document is published because the author believes that people carrying C-PTSD, PTSD, and related conditions should be able to recognize and name what they carry far earlier than he did.

Levy did not recognize his own condition until his forties, after decades of living with patterns he could not see from inside, when a traumatic brain injury accidentally forced him to recognize and address C-PTSD. Levy had to learn to rewire his brain on several levels, and applied what he learned to (re)wiring LLMs. The Somebody test — built to detect displaced Agency in technical language — turned inward and detected displaced Agency in his own cognition. The tool he built for shell programming diagnosed its builder.

In Levy's own words, his public position on this finding:

1. I have C-PTSD.
2. I am in therapy with a Clinical Neuropsychologist.
3. I sought professional guidance for safety through my process.
4. My work had a measurable positive impact on my mental health.
5. Here is what I think about why, from my own experience, not from clinical authority.

If one person reads this and recognizes something in themselves — and takes it to a qualified clinician rather than to an AI — this document has done its intended job.

Invitation

If you are a clinician, researcher, or medical practitioner working in trauma, PTSD, expressive therapies, or AI-assisted therapeutic methods, and this observation interests you, the author welcomes hearing from you.

Levy is interested in any serious engagement — whether as postmortem study of his experience building Grace's Training, or in helping develop clinical methods based on proper trials. The process data from the collaboration exists. The clinician's observations are documented. The practitioner is willing to participate in structured investigation.

What is needed is qualified clinical guidance to determine whether this inadvertent finding is reproducible, generalizable, or an artifact of one practitioner's particular circumstances. The author does not have the standing to answer that question. Clinicians do.

Contact: ike@blackskyresearch.net

*This document accompanies the technical paper *Grace's Training: Transferring Practitioner Judgment into AI Through Plain Text* (Emlin, Levy, Marsh, 2026).*

Full examination of this pathway — its mechanisms, its reproducibility, its implications for practitioner-AI collaboration as a therapeutic vector — is ongoing work Levy continues under clinician care, with interest from additional clinical practitioners.