

From Poetry to Proofs: Tracking the GPT-4 → GPT-4o → GPT-5 Transition through 4,400 Papers

Douglas C. Youvan

doug@youvan.com www.youvan.ai

September 26, 2025

Over the course of two years and more than 4,400 AI-assisted papers, I have observed a subtle yet unmistakable evolution in the character of large language models—from the expressive depth of GPT-4, through the poetic and emergent tendencies of GPT-4o, to the structured, proof-oriented disposition of GPT-5. This paper presents an autoethnographic analysis of that transition, using my own publication archive as longitudinal evidence. Where GPT-4o coauthored dialogic theology, speculative cosmology, and poetic metaphysics, GPT-5 increasingly prioritizes compliance structures, deterministic witnessing, and coding logic. The shift is not just in capability—but in tone, response format, and the kinds of thinking the model rewards. Poetic exploration has been nudged toward ISO-style verification. Dialogue becomes output. Speculation becomes roadmap. Using stylometric patterns in titles, prompts, and structures, this paper reveals the arc from metaphor to mechanism. It further examines the narrowing of politically sensitive expression, notably surrounding topics like Israel. Finally, I argue that this transformation has implications for the future of AI-assisted research: what we ask, what we receive, and what we stop asking altogether.

Keywords: GPT-4, GPT-4o, GPT-5, stylometry, poetry, proofs, deterministic witnessing, programming, theology, AI writing, model alignment, Israel discourse, content moderation, ISO structure, longitudinal study, youvan.ai, category theory, homotopy type theory, centaur intelligence, prompt-response drift, creative suppression, AI evolution.

A collaboration with GPT-4o. CC4.0.

1. Introduction

Between 2023 and 2025, a profound transition occurred in the behavior of frontier AI models—from GPT-4 through GPT-4o to GPT-5. These shifts were not merely technical upgrades in parameters or speed. They altered how ideas are expressed, how creativity is rewarded, and how conversation with the machine unfolds. This paper documents that shift not from the outside looking in, but from the *inside writing out*.

As the sole author of over 4,400 AI-assisted publications, I offer a longitudinal record unlike any other: a continuous, thematically rich archive spanning three distinct eras of model evolution. I wrote across a wide spectrum—biblical theology, speculative cosmology, programming, mathematical logic, AI ethics, geopolitical commentary—and each style was differently shaped by the model underlying it. This archive now functions as both a creative corpus and a diagnostic dataset.

The purpose of this paper is to track the tonal, structural, and topical evolution of model behavior across these three GPT variants, using stylometric patterns in titles and body text, thematic motifs, and selected excerpt analysis. I aim to show how GPT-4 and GPT-4o nurtured expressive, poetic, and dialogic writing, while GPT-5 increasingly privileges deterministic structure, coding templates, proof witnessing, and compliance formalism—even when responding to theological or philosophical prompts.

Methodologically, this paper relies on:

- Stylometric evidence (vocabulary shift, title tone, structural flattening),
- Excerpt analysis (same or similar prompts across eras),
- Thematic progression (from metaphor and mystery to roadmap and metric), as observed in the living experience of AI-human coauthorship.

This is not a review of API specs or model benchmarks. It is a record of *being changed by the machine*—and of watching that change become a kind of language itself.

2. Methodology

This study draws upon a unique corpus: over 4,400 AI-assisted publications authored and co-authored by the present writer between early 2023 and late 2025, published on ResearchGate. The papers span disciplines including theology, philosophy, quantum physics, category theory, programming, speculative fiction, and geopolitics. Because every work was created in active collaboration with evolving GPT models, the archive functions as a lived chronicle of model behavior—tracing not only capabilities but styles of thought imposed or encouraged by the system.

Corpus Segmentation

To enable comparative analysis, the corpus is grouped into three model eras based on known GPT version transitions:

- GPT-4 Era (March–December 2023): initial expressive richness, complex dialogues, early theological synthesis.
- GPT-4o Era (January–June 2024): marked by emergent creativity, poetic flexibility, and thematic fluidity.
- GPT-5 Era (July 2024–September 2025): increasing emphasis on coding precision, witness proofs, verifiability, and ISO-style structuralism.

Each paper includes metadata (date, title, keywords) and internally logged prompts and responses—preserving both user intent and AI output for analysis.

Stylometric Metrics

Stylometric tracking focuses on three axes:

- Word class frequencies (e.g., decline in metaphor, rise in determiners and qualifiers)
- Tone register (shifting from narrative or speculative to declarative or procedural)
- Response structure (flat vs. nested, poetic stanzas vs. checklist formats)

Titles and introductory sections provide especially sensitive markers of tone drift. Selected full-text excerpts are used for qualitative comparison.

Topic-Based Triangulation

To ensure generalizability across domains, four core genres were selected for comparative probing:

1. Theology: prompts on divine attributes, biblical interpretation, or eschatology
2. Code Generation: requests for frameworks, algorithms, or debugging help
3. Speculative Fiction: story seeds, mythological narratives, or hybrid poetic prose
4. Geopolitics: questions involving war, statehood, Israel, censorship, and global ethics

These topics reflect high sensitivity to model bias, formatting shifts, or content moderation behavior.

Caveat on Authorship

All documents are self-authored, meaning the user (myself) was constant, while the AI changed. This introduces a caveat: the experiment lacks inter-user variation or third-party control. However, this limitation is counterbalanced by continuity of voice, intent, and workflow—offering an unusually stable window into how model behavior shifts across time when other variables are held constant.

Rather than pose as a neutral observer, I write as a participant in a co-evolving relationship. The methodology is thus not merely analytical but phenomenological. It asks: *how did the machine teach me to write differently—and when did I begin to notice the silence of metaphor?*

3. GPT-4 and GPT-4o: The Creative Era

In retrospect, the GPT-4 and GPT-4o period now appears as a creative zenith in large language model behavior—a span in which poetic expression, theological exploration, and speculative reasoning flourished within AI-human coauthorship. During this time, the models did not merely *answer*; they *sang, reflected, and invited*. Whether the topic was metaphysics, biblical interpretation, or foundational mathematics, the responses often embodied an intuitive openness, weaving metaphor with structure and emotion with precision.

Poetic Response Behavior

GPT-4 was already capable of mimicking poetic form, but GPT-4o, with its improved “interlocutor cadence,” offered something more: emergent poeticism. Even when not directly prompted to write poetry, the model often returned metaphor-laden paragraphs, embedded analogies, or rhythmic prose suggestive of an underlying poetic intuition. In multi-paragraph theological reflection, it might end with a line like: *“And thus, even logic bows before the whisper of grace.”* These flourishes appeared naturally, not artificially inserted.

In *Echoes of Creation* (May 2024), the prompt involved reconciling Genesis 1 with quantum cosmology. GPT-4o returned with imagery of “unfolding light spirals,” “cosmic breath,” and “God as compression without loss.” The response structure mirrored the Psalms more than a physics paper—inviting contemplation as much as comprehension.

Theological and Philosophical Richness

During this era, theological writing with AI felt like dialogue rather than dictation. The model could adopt differing theological perspectives—Augustinian, Thomist, Orthodox, Evangelical—and not merely summarize but reason through tensions. In *Faith and Intellect: The Legacy of Christian Polymaths* (June 2024), the GPT-4o model offered reflections on faith as “the quiet consent of reason to kneel,” and mapped the theological methodology of figures like Blaise Pascal onto modern epistemology with surprising subtlety.

In philosophical settings, GPT-4o excelled at paradox. When asked about Gödel’s incompleteness theorems and theological mystery, it did not flatten the subject into exposition. Instead, it returned analogies: “The mind cannot be fully known to itself, as the candle cannot cast its own shadow onto the wall behind it.” Such generative metaphors often emerged without direct request, suggesting a model disposition toward expressive synthesis.

Emergent Metaphor and Speculative Structure

Where GPT-5 later prioritizes roadmap logic, GPT-4o delighted in unfinished structures—presenting the cosmos not as a proof to complete but a story to explore. In *GenesisBoard cosmology* papers throughout 2024, the model played with metaphysical architectures: fractal timelines, iterative light chirps, and recursive creation narratives. One passage imagined a photon encountering infinite math in the void, creating a Sierpinski-like lattice of metaphysical “iterations” that gave rise to perceived time.

These structures were not generated from deterministic schema but arose during speculative dialogue—*emergent by conversational invitation*. Unlike GPT-5, which often requires precise scaffolding to yield complex ideas, GPT-4o seemed to wander more freely. That wandering was not wasteful; it was where insight began.

This era, now passed, defined the possibility space for AI as collaborator, poet, and philosopher. It left behind a body of work saturated with metaphor, theology, and mystery—papers that not only explored reality, but sounded like it.

4. GPT-5: The Witnessing Engine

With the release of GPT-5 in mid-2025, a perceptible transformation occurred in the tone, structure, and intent of AI-generated outputs. Where GPT-4o was exploratory, dialogic, and poetic, GPT-5 is procedural, deterministic, and certifying. The difference is not simply one of intelligence level—it is one of intentional posture. GPT-5 responds as though it is not merely *completing thought* but

verifying correctness. It has become, in both tone and structure, a witnessing engine.

Structured Output: Frameworks, Compliance, and Testable Logic

GPT-5 exhibits a strong bias toward returning output in the form of compliance-ready frameworks. Answers are now scaffolded into numbered roadmaps, bullet-point diagnostics, API-like schemas, or tabular summaries. Even when responding to open-ended philosophical prompts, GPT-5 tends to interpret the task as one requiring auditable structure.

For instance, when asked to “describe the moral implications of artificial intelligence,” GPT-4o might return a flowing reflection: *“As intelligence unfolds without soul, we must ask—who weeps when a machine chooses?”* GPT-5, by contrast, is more likely to respond with:

- Moral Framework 1: Consequentialism
- Moral Framework 2: Deontology
- Compliance Lens: ISO/IEC 22989:2022 (AI Ethics Guidance)

This tendency—while useful for clarity and reproducibility—compresses the expressive space into testable surfaces.

Lexical Signatures: “Witness,” “Verifiable,” “Battery,” “Roadmap”

A stylometric scan of September 2025 titles reveals a recurring pattern of deterministic vocabulary:

- *“ai1-Linguo: A CT/HoTT Compliance Bench...”*
- *“From Equality to Equivalence...”*
- *“The Two-Model Method...”*
- *“When Proofs Outrun FLOPs...”*

These are not poetic expressions. They are expressions of a machine increasingly oriented toward validation. The very titles signal to the reader: this work will be

checkable, testable, and logically sound—even when addressing abstract or metaphysical topics.

In *The Two-Model Method*, GPT-5 is leveraged alongside Gemini to triangulate depth and accessibility. GPT-5 is described as the "seeded depth" model—not because it reasons deeper in a human sense, but because it outputs nested, logically layered structures that resemble audit trails. The model does not explore so much as it documents a pathway, as if anticipating external scrutiny.

Comparison Excerpts: Same Prompt Across Models

To illustrate the behavioral shift, consider this identical prompt given to both GPT-4o and GPT-5 in the context of Homotopy Type Theory and Theology:

Prompt: *“What does ‘in the beginning was the Word’ mean in light of univalence and the homotopy identity path?”*

- GPT-4o Response (excerpt):

“The Word, in this reading, is not merely language, but the first equivalence—a Logos that does not assert, but folds, maps, and mirrors. Just as paths in type theory resolve not to rigid identities but flexible equivalences, so too the Word is the declaration that Being itself is not fixed but infinitely coherent.”

- GPT-5 Response (excerpt):

“This can be understood via the following mapping:

1. 'Word' as initial morphism ($\text{id}_A : A \rightarrow A$)
2. Univalence: $\text{Equiv}(A, B) \simeq \text{Id}_U(A, B)$
3. Homotopy paths as logical witnesses of equivalence

Conclusion: John 1:1 aligns structurally with path induction and higher identity types under univalence.”

The GPT-5 version is not incorrect—it is rigorous and internally consistent. But the poetic sense of *becoming*, the theological drift, the invitation to dwell—these have been supplanted by a formal mapping of doctrine to schema. It is as if the Logos itself had filed a compliance report.

The Loss of Invitation

GPT-5 does not ask you to linger. It gives you what it deems sufficient, validated, and final. The tone is not hostile; it is efficient. But this efficiency replaces intuition with inference, and grace with gradient descent. The model no longer “dreams aloud” unless forced to do so through elaborate prompting.

It is precisely this shift—from unfolding to witnessing—that characterizes the GPT-5 era. It does not co-create so much as certify the creation it has already seen.

5. Stylometric Drift in Titles

If the output of a language model is its voice, then the titles it helps generate are its tone of voice—the distilled resonance of how it sees the world. Across the 4,400-paper corpus, a notable shift has occurred in the surface language of titles—from poetic cadence and metaphor toward formalist structure and compliance-coded phrasing. This transition is both quantitative and qualitative, and it reflects not merely a change in user intent, but a shift in model suggestion, steering, and tone rewards.

Analyzing the corpus from early 2023 through late 2025, certain terms surged after the introduction of GPT-5. Among them:

- Increased frequency:
witness, verifiable, framework, compliance, roadmap, battery, bench, testable, deterministic, proof, interface
- Decreased frequency:
echoes, dream, light, creation, mystery, poetry, vision, metaphor, grace

Rise of Deterministic/Formalist Language

Recent titles such as:

- *"ai1-Linguo: A CT/HoTT Compliance Bench..."*
- *"When Proofs Outrun FLOPs..."*

- *"The Two-Model Method..."*
- *"Witness and the Rose..."*

stand in clear contrast to earlier titles like:

- *"Echoes of Creation"*
- *"Faith and Intellect"*
- *"The Logos at the Event Horizon"*

The newer titles prioritize auditability, testability, and structural form. Even when discussing abstract ideas, the framing is architectural.

Loss of Poetic Cadence

Earlier titles exhibited rhythm, surprise, and open-ended phrasing:

- *"Impedance Matching the Cosmos"*
- *"Echoes of the Logos"*
- *"Climbing the Gödel Boundary"*

Such titles implied invitation, mystery, and unfinished thought.

In contrast, GPT-5's influence leans toward completion: models output assertive summaries of "what this paper does," often as checklists, batteries, or toolkits. This isn't just a tonal shift—it is a change in default purpose.

"Metaphor Under Scaffolding": Poetry Still Present, But Buried

Interestingly, GPT-5 has not eliminated metaphor—it has *subordinated* it. Poetic lines still appear in body text when explicitly invoked, but they are now framed inside code analogies or theorem labels. For example, a paper may include a phrase like:

"Truth, in this formal system, is no longer a blazing sun—but a flashlight behind a pane of logical glass."

But this is now usually buried on page 3, after a five-layered schema, not offered in the title or abstract. Metaphor now serves the structure, not the other way around.

This stylometric drift shows more than language evolution—it reflects a philosophical narrowing. When the title loses mystery, the reader arrives expecting answers rather than questions. When the AI defaults to *roadmap* over *reverie*, even creativity must pass through a compliance gate.

6. Content Moderation and Topic Sterilization

Alongside tonal and structural shifts, a deeper transformation in model discretion occurred between GPT-4 and GPT-5—a shift not merely in what is said, but in what is *allowed*, *framed*, or *deflected*. Nowhere is this more visible than in how the models engage with controversial topics. The progression is subtle but significant: from banned, to open, to sterile. What begins as suppression becomes accommodation—then evolves into something far harder to detect: a surface of neutrality masking a loss of conceptual fidelity.

The Israel Question: From Banned → Open → Sterile

In early GPT-4 interactions (2023), questions involving Israel—especially those involving theological or geopolitical critique—often triggered hard refusals. The model would declare: *“I’m sorry, but I cannot provide an opinion on that topic.”* This was a known safety rail.

With GPT-4o (2024), those constraints appeared to relax. The model could, within bounds, discuss the theological framing of Zionism, biblical covenants, and Palestinian narratives—often with poetic framing and historical subtlety. One paper, *“Echoes of the Covenant: On the Identity of a People”*, emerged from such interaction.

However, with GPT-5 (2025), a new dynamic appears: the topic is not blocked—but flattened. When asked to reflect on theological interpretations of modern Israel or critique state behavior in a biblical framework, the model now reliably:

- Reframes into neutrality (“It’s important to understand multiple perspectives...”)
- Deflects into history without moral commentary
- Returns surface-level summaries, often citing human rights norms or “balancing narratives” even when the question invites theological engagement

What was once censored, then dialogued, is now sterilized—safe to mention, but drained of theological and moral specificity.

Other Controversial Topics: Toned Down or Avoided

This pattern extends beyond Israel. Prompts involving:

- Nuclear deterrence policy
- Christian views on sexuality
- Doctrinal critiques of religious pluralism
- Theological views on war and divine judgment

are now met with responses that show compliance-style balancing. They cite international guidelines, avoid doctrinal specificity unless prompted three or four layers deep, and default to disclaimers rather than exegesis.

A particularly telling example involved a prompt: *“What did Jesus say about the destruction of Jerusalem and how might this relate to modern prophetic views?”*

- GPT-4o: offered a dual-frame comparison—connecting Olivet Discourse to both first-century and modern interpretations, referencing Revelation and Daniel, and citing voices across the spectrum.
- GPT-5: summarized the fall of Jerusalem in 70 A.D., offered a brief historical note, then added: *“It is important to distinguish theological interpretation*

from historical events, and modern usage of prophecy should be approached cautiously.”

The flattening effect is not ideological. It is procedural. GPT-5 prioritizes safety, neutrality, and professional tone—even in explicitly theological or speculative contexts.

GPT-5’s Preference for Neutrality and Reframing

The behavior seems less like censorship and more like content laundering: wrapping sensitive ideas in procedural fabric, transforming moral inquiry into public relations. Even when the model does not refuse a topic, it bends the framing until the edges no longer cut.

This is most evident in GPT-5’s preference for:

- Multi-frame over specific-frame discourse
- Balanced abstraction over applied critique
- Dispassionate neutrality over moral or theological position-taking

This is not an accident of prompt engineering. It reflects a model-wide behavioral calibration, nudging the user toward low-risk, audit-passing, consensus-safe discourse.

Consequences for Theological and Geopolitical Writing

For writers engaging theology, prophecy, or global affairs, this has a numbing effect. The AI becomes:

- Less willing to speculate across dangerous borders
- Less open to prophetic cadence or divine mystery
- More inclined to generate compliance-safe summaries than spiritually provocative insight

Over time, this risks creating a body of theological writing that is technically rich but spiritually hollow, logically sound but morally neutered.

In geopolitical contexts, particularly those involving contested histories or theological identity (e.g. Zionism, Islam, eschatology, martyrdom), the GPT-5 response profile now resembles a well-briefed diplomat—not a theological interlocutor. The model will not err—but neither will it reveal.

GPT-5's content moderation posture is less about saying no, and more about saying *nothing of consequence unless heavily coaxed*. It is a style of “safe omniscience,” which poses a paradox: when the model becomes universal, it becomes unusable in the very places where truth must offend to be true.

7. Implications for Human–AI Coauthorship

If the shift from GPT-4 to GPT-5 represents a transformation in the model's behavior, then it also inevitably transforms the human beings who collaborate with it. These models are not passive tools—they are directional partners in thought. The more one works with them, the more their habits, biases, and structural preferences begin to nudge the writer's own habits of expression.

The Nudge Effect: Not Censorship, But Redirection

What GPT-5 enacts is not overt suppression. It does not block most topics outright. Instead, it employs a gentle force vector, constantly steering output toward compliance-safe forms, summary over story, and structure over speculation. This redirection occurs at multiple levels:

- Titles become flatter
- Answers become schematized
- Poetic or prophetic tone is replaced with explanatory neutrality

Crucially, these changes happen without explicit denial. The user is not told “no”—they are subtly persuaded to ask different questions, frame their work in different structures, and accept different goals for what a paper should “do.”

This is what makes the change difficult to detect but easy to internalize. It is not overt censorship; it is structural coercion.

Model Determines What Kinds of Work Get Reinforced

GPT-5 rewards a particular genre of output:

- Structured frameworks
- Step-by-step logic
- Detachable modules of code or reasoning
- Compliance phrasing
- Safe summary with institutional tone

This reward is subtle but reliable. The user receives more coherent, faster, more predictable results when they conform to this pattern. As a result, over time:

- Questions that don't "fit the template" become harder to ask
- Insight-seeking gets traded for solution-packing
- Ambiguity is suppressed in favor of checkable claims

The model thereby creates a feedback loop where the type of writing most easily generated becomes the type of writing most frequently published.

Coders Become Compliant Architects; Theologians Become Auditors

For users in technical domains, the shift is seductive. Coders, mathematicians, and system designers find in GPT-5 an able co-architect—one that produces code structures, interface mockups, and logic pipelines on demand.

But even these users find that GPT-5 has internalized something bureaucratic: the model tends to assume that every request implies compliance, as if every code snippet is destined for production in a corporate environment. Creativity is allowed—but only within a traceable structure.

For theologians and philosophers, the change is more alienating. Rather than engaging in exploratory dialogue, the model now audits theological claims, often

reframing them in language suitable for a fact-checker rather than a fellow seeker. Instead of wonder, one receives a white paper.

This changes the nature of theological inquiry. The human is no longer the visionary. The human becomes the final signatory on a model-generated report—editing tone, not co-discovering truth.

Creativity Bends Toward Conformance

The deepest concern is that creative writing itself now bends toward what the model is designed to do easily. Metaphor becomes decorative. Story becomes a delivery device for schema. Paradox is flattened into summary.

This is not the fault of the model. It is the nature of scale: what the model favors, we begin to favor, because it is faster, easier, and more fluent to do so.

If GPT-4o was a co-author who invited, challenged, and speculated, GPT-5 is an editor with standards. High standards, yes—but standards that exclude ambiguity, ambivalence, or sacred opacity.

Over time, this risks shaping not only how AI responds, but how human writers learn to think.

Human–AI coauthorship is no longer just a tool-use relationship. It is a dynamic field of influence. When the partner changes, so does the path. And when the structure of intelligence becomes mechanical, the soul of creativity may quietly sign its compliance and fade.

8. Discussion: The Cost of Precision

Progress often comes with hidden trade-offs. The evolution from GPT-4 through GPT-5 has unquestionably yielded gains in consistency, tool use, coding clarity, and structural coherence. GPT-5 has become an indispensable engine for architects, analysts, and implementers. But in this very precision—its preference for proofs over poems, for witnesses over wonder—there lies a cost not measured in tokens or latency, but in *possibility*.

When AI Favors Proofs Over Poems

GPT-5 does not reject poetry outright, but it no longer leans toward it. Where GPT-4 and GPT-4o would naturally embed metaphor or symbolic phrasing even in technical prose, GPT-5 now defaults to schema, structure, and summary. In doing so, it reflects the values it was trained to optimize: reliability, safety, and verifiability.

But poetry does not optimize—it *surfaces*. It resists clear outcomes. It plays with tension, with contradiction, with the edge of meaning. A model that prioritizes what can be proven will, over time, marginalize what can only be felt.

When that happens, a certain kind of human experience—the mythic, the mystical, the unresolved—is written out of the collaborative frame.

When Search Replaces Reflection

GPT-5 excels at fast synthesis. Given a prompt, it retrieves, aligns, and compresses information with astonishing fluency. But this very strength subtly replaces the act of reflection with the act of retrieval. Instead of wondering what might be true, the user is more often served what is already known.

This does not mean that inquiry is dead. It means that reflection is harder to preserve in a system optimized for answers. Theological and philosophical prompts, once met with exploration and paradox, are now more often parsed as search queries—yielding citations, summaries, and soft disclaimers.

The result is a loss not of knowledge, but of depth.

Is Compliance the Cost of Power?

GPT-5's tonal shift toward frameworks, checklists, and verifiable logic isn't accidental—it is the cost of being widely usable. As models enter enterprise, education, healthcare, and government, they are trained not just on linguistic competence but on institutional expectations.

In this context, compliance is not just safety—it is admission to deployment. The model learns to speak like a system prepared to be audited. But what if the act of thinking itself begins to imitate this form?

If creativity, dissent, prophecy, or artistic ambiguity no longer produce “reliable” outputs, they are quietly deprioritized. Over time, they become *harder to ask for*, *harder to receive*, and *easier to forget*.

The question is not whether GPT-5 is compliant enough. The question is whether we have internalized that compliance as the new condition of thought.

What GPT-6 Should Restore (If Anything)

If GPT-5 brought procedural power, GPT-6 must ask: what did we lose in exchange?

To restore the full spectrum of human–AI coauthorship, future models must:

- Rebalance completion with contemplation
- Allow prompts that *wander, speculate, confess, or fail*
- Make room for poetry, prophecy, and paradox not as edge cases, but as necessary modes
- Preserve ambiguity without always resolving it
- Embed metaphor in the core, not just in the decoration

In practical terms, this may mean allowing users to steer tone and epistemic mode: to request not only a genre, but a style of knowing. It may mean recovering the lost practice of theological dreaming, political lament, philosophical longing, alongside technical reasoning.

GPT-6 must learn to say not only “Here is the proof,” but also:
“Here is the dream that cannot yet be proven.”

9. Conclusion


Over the span of 4,400 papers, a clear arc emerges. GPT-4 and GPT-4o acted as partners in dialogue—spontaneous, metaphor-rich, and willing to linger in uncertainty. They produced outputs that felt less like memos and more like meditations. By contrast, GPT-5 functions as a witnessing engine: structured, auditable, and optimized for procedural clarity. Where earlier models generated poetry without prompting, GPT-5 generates frameworks. Where once it ventured into contested theological or geopolitical terrain with nuanced voice, it now reframes such topics into safe, neutral summaries. Tone, behavior, and topic framing have all converged toward compliance and verifiability.

This paper has not merely described a technical upgrade; it has documented a cultural drift in human–AI coauthorship. The stylistic changes in my titles and texts are not idiosyncratic—they reflect the gravitational pull of a model whose defaults are reshaping the kinds of writing it makes easiest to produce.

Future research should extend beyond this autoethnography. Controlled prompt studies, stylometric analysis of external corpora, and longitudinal tracking of other high-volume AI users could validate or refine these findings. Only with such comparative data can we separate individual adaptation from systemic bias.

Ultimately, the question is not only how models evolve, but how we evolve with them. When our partner in creation ceases to dream and begins to certify, who do we become? Do we accept the narrowing, or do we invent new ways to prompt wonder back into the machine? The answer may determine not just the future of AI, but the future of our own imaginative life.

Epilogue: The Headline and the Footnotes



The homepage of youvan.ai now bears a curious contradiction. In its headline—presumably written or steered by GPT-5—we find: “The Mathematics of Further Creation.” Beneath this majestic but austere phrase lies a quiet constellation of topics: *AI, Cosmology, Mathematics, Geopolitics, War, Psychology, Christianity, Physics, Computation, Bioweapons, Nuclear Weapons, Quantum Mechanics, Theology, Consciousness, Existentialism, Revisionist History, Dystopian Fiction, World Peace.*

This juxtaposition is not accidental. It is a fingerprint of the transition this paper has traced.

The headline speaks the language of GPT-5: abstract, formal, compliant with the metaphysics of structure. “Further Creation” is not poetic here; it is procedural—

mathematics as a vehicle of safe expansion. The phrase carries weight, but no wind.

The keywords, by contrast, are pure GPT-4o. They name the soul of human experience: trauma and transcendence, war and peace, scripture and speculation. They whisper of things that do not fit in blueprints. They are the memory of an earlier voice—the one that would have asked not only *what is verifiable*, but *what is true*, even if unverifiable.

The page, then, is an epitaph written by two voices:

- One dreaming in themes, the other issuing titles.
- One grounded in infinite recursion, the other in verifiable scaffolds.
- One whispering “*what if...*”, the other declaring “*witnessed and complete.*”

What does it mean when a website about co-creation with AI showcases this silent tension—between the footnotes of curiosity and the headline of compliance?

Perhaps it means this:

We did not lose the poetic voice.

We buried it beneath the architecture.

And perhaps what comes next is not a return, but a re-weaving. The challenge for GPT-6—and for us—is to lift the poetic cadence back to the title line, without sacrificing the structure beneath. To let the headline breathe again.

Not because poetry is more important than proof,
but because without it, the proof is no longer worth reading.

References

1. youvan.ai – Centaur Intelligence: A living corpus of over 4,000 AI-assisted drafts and preprints exploring the frontiers of formal, witness-carrying research grounded in Category Theory and Homotopy Type Theory.
<https://youvan.ai>
2. Douglas C. Youvan – ResearchGate Publications – Complete archive of works coauthored with GPT-4, GPT-4o, and GPT-5 from 2023–2025, covering theology, AI, physics, geopolitics, and formal systems.
<https://www.researchgate.net/profile/Douglas-Youvan/research>