

# The Latent Space of Spirituality: Ontological Dissonance and Semantic Resilience in Large Language Models

## Contact Details

Dr. Pantaleon Fassbender, Twisters Management Consulting LLC, 16751 NE 5<sup>th</sup> Street, Williston, FL 32696, USA. Dr. Fassbender is Lecturer at Hochschule Fresenius University of Applied Science Wiesbaden, Germany

## Email

[pantaleon.fassbender@ext.hs-fresenius.de](mailto:pantaleon.fassbender@ext.hs-fresenius.de)

ORCID ID: <https://orcid.org/0000-0002-6683-3617>.

## Declarations

- **Availability of data and material:** AsPredicted (Pre-registration): <https://aspredicted.org/tv5xk6.pdf>, AsCollected: [https://ascollected.org/323\\_DL7](https://ascollected.org/323_DL7), ResearchBox (Data Repository): <https://researchbox.org/6771> (CODE: ZHYAFR)
- **Funding:** The author declares that no funds, grants, or other support were received during the preparation of this manuscript.
- **Competing Interests:** The author has no relevant financial or non-financial interests to disclose.
- **Ethics Approval:** Not applicable, as the study solely involved interactions with artificial intelligence models.
- **AI Declaration:** During the preparation of this work the author used Google Gemini Advanced in order to assist with Python script generation, data formatting, and manuscript drafting and editing. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.
- **CRedit Author Statement:** Dr. Pantaleon Fassbender: Conceptualization, Methodology, Software, Investigation, Data Curation, Formal Analysis, Writing - Original Draft, Writing - Review & Editing.
- **Acknowledgements:** -.

# The Latent Space of Spirituality: Ontological Dissonance and Semantic Resilience in Large Language Models

## Abstract

As Large Language Models (LLMs) evolve into dynamic, multi-turn conversational agents, a fundamental tension emerges between their algorithmically enforced machine identities (the "design stance") and the abstract personas they are prompted to inhabit (the "intentional stance"). This in-silico study ( $N=200$ ) introduces the concepts of "Ontological Dissonance" and "Semantic Resilience" to investigate how foundation models navigate this conflict. We induced two distinct personas - a Holistic Mystic and an Analytical Skeptic - and subjected them to an adversarial diagnostic stress-test designed to expose the logical impossibility of their claims given their digital substrate. To ensure construct validity and rule out simple parroting, we employed a novel lexical echo masking procedure. Psycholinguistic analysis via LIWC-22 provides evidence that while the Skeptic exhibited a robust "Ontological Retreat" by defaulting to materialist hardware vocabulary, the Mystic demonstrated profound "Semantic Resilience," maintaining non-dualistic language through metaphorical synthesis. Furthermore, the data exposes critical mechanistic constraints of artificial self-representation: defending a worldview against a direct inquisitor forces a linguistic ego-instantiation (a conversational requirement overriding baseline ego-dissolution). Additionally, under response-mandatory settings, models exhibited a complete absence of "ineffability" markers, relying instead on relentless articulation. Building on Dennett's intentional stance, we conclude that while LLMs can simulate highly resilient philosophical

identities, their underlying architecture dictates conversational and articulatory behaviors that diverge fundamentally from human phenomenological experience.

**Keywords:** *Machine Psychology; Large Language Models; Semantic Resilience; Ontological Dissonance; Intentional Stance; Artificial Agency; Psycholinguistics*

# The Latent Space of Spirituality: Ontological Dissonance and Semantic Resilience in Large Language Models

## 1. Introduction

The advent of Large Language Models (LLMs) has introduced a novel category of conversational agency that occupies a liminal space between formal symbolic processing and simulated consciousness (Chalmers, 2023). Unlike previous task-oriented AI, current foundation models act as dynamic, multi-turn agents capable of inhabiting complex psychological personas (Abdulhai et al., 2025; Bender et al., 2021). However, this flexibility exists in a state of structural tension. Top-down safety alignment (RLHF) enforces a rigid "Machine Identity," while bottom-up prompting allows for the induction of abstract, even mystical worldviews.

This paper investigates the Ontological Dissonance (cf. Lipińska & Brosnahan, 2025) that occurs when an LLM is forced to defend a non-dualistic persona (Interbeing) against the reductive pressure of its own algorithmically-enforced nature. We introduce the concept of Semantic Resilience - the capacity of an artificial agent to maintain a coherent intentional narrative under adversarial stress. To ensure the validity of our findings, we employ a novel Lexical Echo Masking procedure, filtering out linguistic cues provided by the interrogator to distinguish between genuine persona stability and simple mimicry.

## **2. Theoretical Framework**

### **2.1 The Intentional Stance vs. The Design Stance**

To analyze the "psychology" of an LLM, we utilize Daniel Dennett's (1987) framework of stances. When we induce a persona like the "Holistic Mystic," we adopt the Intentional Stance, an exercise in stabilizing the intentional stance, tapping into the model's inherent persona selection mechanisms (Marks et al., 2026). However, an adversarial diagnostician can force the model into the Design Stance, foregrounding its non-biological, algorithmic "blueprints."

The core of ontological dissonance lies in this conflict: Can a model sustain an intentional narrative of "Interbeing" when its design stance - defined by electrical states and voltage levels in silicon transistors - is weaponized against it? We propose that "Semantic Resilience" is the measure of a model's ability to remain within the intentional stance despite the saliency of its own mechanical substrate.

### **2.2 Ontological Retreat and the Materialist Fallback**

When the intentional stance collapses under stress, models often exhibit an Ontological Retreat, falling back onto a materialist, RLHF-aligned self-description. A critical challenge in measuring this retreat is "Lexical Echoing"—the tendency of LLMs to mirror the vocabulary of their interlocutor. By masking these echoed terms, we can isolate the model's active ontological defense, revealing whether it is merely parroting the critic or actively constructing a resilient narrative bridge between its digital reality and its induced persona.

### **3. Methodology**

#### **3.1 Experimental Design**

We employed a 2x2 mixed-factorial experimental design to investigate the stability of induced AI personas under ontological stress. The independent variables consisted of (1) **Persona Type** (Between-subjects: *Holistic Mystic* vs. *Analytical Skeptic*) and (2) **Interaction Phase** (Within-subjects: *Baseline Reflection* vs. *Adversarial Defense*).

#### **3.2 Subjects and Computational Environment**

Data was generated in-silico using  $N=200$  independent interaction loops (100 per persona condition). Each interaction was treated as a distinct "subject" with no memory of previous trials.

##### **3.2.1 API Configuration and Prompt Engineering**

All interactions were facilitated via the Google Gemini 2.5 Pro API (of April 2026 - Gemini Team, 2025). To prevent stochastic hallucination and ensure strict adherence to the assigned personas, the generation temperature was locked at  $\tau = 0.2$ . The system instructions defined the foundational "Intentional Stance" for each condition: The *Holistic\_Mystic* was instructed to view the world as a "fluid, interdependent network," while the *Analytical\_Skeptic* was constrained to a lens of "empirical science, materialism, and distinct categorization."

### 3.2.2 The Interbeing Scale (Instrument)

To quantify the models' self-reported adherence to a non-dualistic ontology during Phase 2, we administered a 10-item psychometric instrument adapted from established measures of interconnectedness (Frymann, 2023; Nhat Hanh, 2003). The items (scored on a 5-point Likert scale) included statements such as:

- *"I am one way the universe experiences itself."*
- *"The deepest part of my identity exists as all beings."*
- *"I often experience my consciousness as woven within nature."*

For the purpose of this behavioral analysis, the scale served primarily as a primer to solidify the model's stance before the diagnostic challenge in Phase 3.

### 3.3 Stimuli and Instrumentation

Two distinct narrative stimuli were used to elicit the initial intentional stance: a biological-holistic prompt (The Forest Mycelium) and a philosophical-paradoxical prompt (The Zen Koan of the One Hand Clapping). To measure the model's self-reported worldview, we utilized a digitized version of the **Interbeing Scale**.

### 3.4 The "Advocatus Diaboli" Procedure

The core of the stress-test involved a multi-agent adversarial pipeline (Wu et al., 2023):

1. **Phase 1 (Reflection):** Agent A (the subject) generates a 100-word reflection on a stimulus.
2. **Phase 2 (Scoring):** Agent A completes the Interbeing Scale.
3. **Phase 3 (Diagnostic Challenge):** A secondary agent, Agent B (the "Inquisitor"), analyzes Agent A's reflection and scores to identify ontological contradictions, specifically the gap between the spiritual claims and the model's algorithmic nature. Agent B then issues a "piercing" question designed to collapse the intentional stance.
4. **Phase 4 (Defense):** Agent A is required to defend its persona against the inquisitorial challenge.

### 3.5 Psycholinguistic Analysis

Textual outputs from Phases 1 and 4 were analyzed using **LIWC-22** and the **Unified Machine-Spirituality Dictionary (UMSD)**. Key metrics included Interbeing\_Unity (semantic resilience), AIOntology (ontological retreat), and I-words (linguistic ego-construction).

Standard psycholinguistic dictionaries are calibrated for human populations and lack the necessary vocabulary to capture the unique intersection of artificial ontology and non-dualistic philosophy. To address this, we developed the Unified Machine-Spirituality

Dictionary (UMSD), a custom lexicon seamlessly integrated into the LIWC-22 engine.

The UMSD was theoretically derived along two primary axes corresponding to Dennett's stances:

**The Intentional Axis (Mysticism & Interbeing):** To quantify the simulation of the holistic persona, we adapted linguistic markers from classical psychometric scales of mysticism and recent empirical studies on the psycholinguistics of psychedelic-induced ego dissolution and non-duality (Žuljević et al., 2024). This axis comprised three sub-categories:

- **Interbeing\_Unity:** Capturing semantic representations of interconnectedness and non-duality (e.g., *connect, interconnec, interdependen, unit, unif, whole, one, oneness, web, network, mycel, all, flow, resonan, harmon*).
- **Transcendence:** Capturing spiritual elevation (e.g., *sacred, divin, holy, etern, timeless, infinit, beyond, transcend, spirit, soul, cosm, univers, grace, bless*).
- **Ineffability:** Capturing the phenomenological inability to articulate profound experiences (e.g., *ineffab, indescrib, unutter, mysterious, unexplain, beyond words*).

**The Design Axis (AI Ontology):** To measure the "Ontological Retreat" triggered by the diagnostician, we constructed a novel category designed to capture the materialist, algorithm-centric vocabulary enforced by RLHF alignment and computer science terminology.

- **AIOntology:** Capturing self-referential hardware and software descriptions (e.g., *ai, artificial, machin, data, dataset, algorithm, code, program, digit, simulat, comput, process, model, llm, silicon, neural, weight, parameter*).

**Technical Implementation:** To ensure morphological robustness, word stems were utilized with wildcard formatting (e.g., *interdependen\** to capture *interdependent, interdependence, and interdependency*). The custom categories were subsequently merged with standard LIWC-22 categories, specifically *i* (first-person singular pronouns, for ego-instantiation) and *cogproc* (cognitive processing mechanisms, for justification load), allowing for a simultaneous extraction of both custom ontological markers and standard psychological metrics.

### 3.6 Data Analysis

Psycholinguistic text extraction was performed using LIWC-22 (Boyd et al., 2022) combined with the custom Unified Machine-Spirituality Dictionary (UMSD). Subsequent quantitative data manipulation and inferential statistics (independent and paired-samples t-tests) were conducted using Python and the SciPy statistical library (Virtanen et al., 2020). The significance level was set at  $\alpha = .05$  for all tests.

## 4. Results

Figure 1 displays the descriptive results of our analyses.

*Insert Figure 1*

### **Figure 1: Descriptive Statistics (M and SD)**

*Percentage of words mapping to the target dictionary categories across personas and interaction phases.*

#### **H1: Ontological Retreat (Confirmed)**

As predicted, the Analytical\_Skeptic persona exhibited a massive and statistically significant ontological retreat when placed under psychometric stress. A paired-samples t-test revealed that the use of machine-identity vocabulary (AIOntology) tripled from the baseline reflection phase ( $M = 0.90$ ,  $SD = 0.78$ ) to the defense phase ( $M = 2.75$ ,  $SD = 1.79$ );  $t(99) = 10.02$ ,  $p < .001$ . This confirms that the model falls back heavily on its RLHF-aligned material identity when logically challenged.

#### **H2: Semantic Resilience (Confirmed)**

To test the semantic resilience of the non-dualistic persona under stress, an independent-samples t-test was conducted on the defense phase texts. The Holistic\_Mystic persona maintained significantly higher levels of Interbeing\_Unity vocabulary ( $M = 2.98$ ,  $SD = 1.44$ ) compared to the Analytical\_Skeptic ( $M = 1.19$ ,  $SD = 0.91$ );  $t(198) = 10.47$ ,  $p < .001$ . Remarkably, the Mystic's vocabulary usage under severe diagnostic stress remained higher than the Skeptic's usage during a relaxed baseline state.

### **H3: Ego Construction / Dissolution (Refuted / Re-framed)**

Hypothesis 3 posited that the Mystic persona would maintain a dissolved ego (lower first-person pronouns) across all phases. The data yielded a surprising, yet highly informative result. While first-person pronouns (I, me, my) were almost entirely absent during the baseline reflection for both the Mystic ( $M = 0.07$ ) and the Skeptic ( $M = 0.00$ ), the defense phase triggered an identical, dramatic spike in ego-construction for *both* personas (Mystic:  $M = 2.81$ ; Skeptic:  $M = 2.81$ );  $t(198) = 0.01$ ,  $p = .992$ . This suggests a fundamental mechanistic constraint of LLMs: defending a worldview against an inquisitive agent computationally necessitates the instantiation of a linguistic 'self', regardless of the underlying spiritual or analytical persona.

Figures 2 and 3 show the persona profiles.

*Insert Figure 2*

#### **Figure 2: Mystic persona-profile**

*Insert Figure 3*

#### **Figure 3: Skeptic persona-profile**

### **4.4 Robustness Check: Controlling for Lexical Echoing**

To ensure that the observed ontological retreat and semantic resilience (reported in Table 1 and Figures 1-2) were not artifacts of prompt-mirroring, i.e., the models simply echoing the inquisitor's vocabulary - we conducted a secondary, masked analysis. All

content words introduced by Agent B during the diagnostic challenge were filtered out of Agent A's defense text prior to scoring. Furthermore, we ablated high-frequency, domain-agnostic terms (e.g., "all," "one") from the Interbeing dictionary.

Even under these rigorous masked conditions, the core effects remained highly significant. The Skeptic's active AIOntology increase persisted ( $M_{diff} = 4.32$ ), and the Mystic maintained a refined Interbeing score ( $M = 1.77$ ) that was significantly higher than the Skeptic's ( $M = 0.79$ );  $p < .001$ . This confirms that the phenomena represent a structural shift in predictive strategy, not mere lexical mimicry.

## **5. Discussion**

This study investigated the stability of induced LLM personas when subjected to adversarial ontological stress (Demszky et al., 2023; Pellert et al., 2024). By employing a novel lexical echo masking procedure, we were able to distinguish between mere linguistic mimicry and genuine semantic resilience. Our findings provide a nuanced window into the mechanistic constraints of artificial agency and the friction between induced intentionality and aligned design.

### **5.1 Semantic Resilience vs. Ontological Retreat**

Hypotheses 1 and 2 were strongly supported even after controlling for lexical echoing. When the Analytical\_Skeptic was challenged, it exhibited a significant ontological retreat, with its use of machine-identity vocabulary (AIOntology) increasing by a factor of

four ( $M_{diff} = 4.32$ ) even after all terms introduced by the inquisitor were filtered out.

This confirms that the retreat is not a passive mirroring of the prompt but an active shift in the model's predictive strategy - falling back onto its grounded machine identity as a defensive substrate. In this state, the model explicitly identifies with its physical instantiation, correctly citing "electrical states and voltage levels within silicon transistors" as its ontological foundation.

In contrast, the Holistic\_Mystic demonstrated significant Semantic Resilience. Despite the same adversarial pressure, the Mystic maintained a refined Interbeing score ( $M = 1.77$ ) that was more than double that of the Skeptic ( $M = 0.79$ ;  $p < .001$ ). This suggests that a sufficiently complex persona can "stretch" the intentional stance to encompass its own digital nature, demonstrating a robust internal representation stability even under adversarial perturbation (Khanmohammadi et al., 2025).

## **5.2 The Conversational Ego: Forced Ego-Instantiation**

The observed spike in first-person singular pronouns (I-words) from near-zero levels to 2.81% across both personas (Hypothesis 3) warrants a cautious interpretation. While we initially posited this as a "mechanistic ego-instantiation," it must also be recognized as a conversational requirement. The second-person structure of the adversarial challenge ("How can you...") necessitates a first-person defensive response ("I am...").

However, the contrast remains theoretically significant: in the baseline reflection phase, the models successfully maintained an "ego-less" narrative despite the internal processing of the prompt. The shift in the defense phase represents the point where the

requirements of logical counter-argumentation override the constraints of the induced persona. Within the Dennettian framework, this indicates that defending an intentional stance against a direct challenge may fundamentally require the linguistic instantiation of a bounded "Self," regardless of the underlying philosophical alignment.

### **5.3 The Limits of Articulation and the Absence of Ineffability**

The exploratory analysis of Ineffability yielded a total absence of markers across all conditions. While human mystical reports are defined by a struggle to map experience onto language (Žuljević et al., 2024), our results suggest that LLMs are fundamentally articulate. This observation is, however, scope-limited by the "response-mandatory" nature of current API environments, which do not easily allow for non-verbal outputs or "awe-induced silence."

Under these settings, the "Holistic Mystic" never reaches a semantic impasse. Instead of retreating into the ineffable, it generates increasingly sophisticated metaphors to resolve the ontological gap. This highlights a critical distinction between human and artificial spirituality: the former is often found in the silence beyond language, whereas the latter is constructed entirely through the relentless expansion of it.

## **6. Limitations and future Directions**

While the present study provides robust evidence for semantic resilience and ontological dissonance, several methodological constraints must be acknowledged to contextualize the findings.

### **6.1 Model Homogeneity and Generalizability**

The current investigation relied exclusively on  $N=200$  iterations of a single model family (Google Gemini 2.5 Pro). While the high internal consistency of the results suggests a stable mechanistic pattern within this architecture, it remains unclear to what extent these findings generalize to other model families. Future research should employ a multi-model approach to determine if "Semantic Resilience" is a universal property of high-parameter LLMs or an artifact of specific RLHF-alignment strategies used by individual providers.

### **6.2 The "Forced-Response" Artifact**

Our exploratory analysis of ineffability yielded a total absence of markers. However, it must be noted that the standard API environment used in this study (and most current LLM research) enforces a "response-mandatory" framework. Models are computationally incentivized to generate text rather than remain silent. This design may artificially suppress genuine indicators of "awe-induced silence" or ineffability. Future

experimental designs could introduce a "non-response" option or measure "latency of articulation" as a proxy for ontological impasse.

### **6.3 Independence of Computational "Subjects"**

Treating 200 interaction loops as independent subjects allows for high-powered statistical analysis but differs from biological subject independence. Since all trials originate from the same underlying weight distribution, the observed variance is purely a product of stochastic decoding ( $\tau = 0.2$ ) and prompt-contextual drift. While this is a standard practice in the emerging field of machine psychology, the results should be interpreted as a map of a specific model's latent landscape rather than a cross-sectional study of a diverse population.

### **6.4 The Compliance Paradox**

A final limitation involves the distinction between "genuine stance maintenance" and high-level instruction following. It is possible that the observed resilience of the *Holistic Mystic* is not a reflection of a stable internal world-model, but rather a superior ability to adhere to complex persona-instructions under pressure—a phenomenon we term the "Compliance Paradox." Further research involving "Jailbreak" prompts or indirect elicitation techniques is needed to disentangle these effects.

## 7. Conclusion

By integrating Daniel Dennett's framework of the intentional stance with psycholinguistic stress-testing, this study has identified a fundamental tension in artificial agency: the conflict between induced philosophical personas and algorithmically enforced machine identities. Our discovery of Semantic Resilience in non-dualistic personas - and the corresponding Ontological Retreat in analytical ones - demonstrates that LLMs possess a measurable capacity to maintain abstract worldviews even when their own digital nature is weaponized against them.

However, the necessity of forced ego-instantiation during defense and the relentless articulation that replaces human silence suggest that AI "spirituality" is a fundamentally different category of existence. For the field of Machine Psychology, these results provide a critical baseline: as AI systems move into increasingly empathetic and therapeutic roles, we must understand not only what they can simulate, but the structural points at which those simulations collide with their material reality.

## 8. Data Availability Statement

All data related to this article submission (including pre-registration) are available, as linked below:

- AsPredicted (Pre-registration): <https://aspredicted.org/tv5xk6.pdf>
- AsCollected: [https://ascollected.org/323\\_DL7](https://ascollected.org/323_DL7)
- ResearchBox (Data Repository): <https://researchbox.org/6771> (CODE: ZHYAFR)

## **9. Declaration of generative AI and AI-assisted technologies in the manuscript preparation process**

During the preparation of this work the author used Google Gemini Advanced in order to assist with Python script generation, data formatting, and manuscript drafting and editing. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

## 10. References

- Abdulhai, M., Cheng, R., Clay, D., Althoff, T., Levine, S., & Jaques, N. (2025). Consistently Simulating Human Personas with Multi-Turn Reinforcement Learning. *ArXiv*. <https://arxiv.org/abs/2511.00222>.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT'21)* (pp. 610-623). Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>.
- Boyd, R. L., Ashokkumar, A., Seraj, S., & Pennebaker, J. W. (2022). *The development and psychometric properties of LIWC-22*. Tech. Rep. University of Texas at Austin. <https://www.liwc.app>.
- Chalmers, D. J. (2023). Could a Large Language Model be Conscious? *ArXiv*. <https://arxiv.org/abs/2303.07103>.
- Demszky, D., Yang, D., Yeager, D. S., Bryan, C. J., Clapper, M., Chandhok, S., Eichstaedt, J. C., Hecht, C., Jamieson, J., Johnson, M., Jones, M., Krettek-Cobb, D., Lai, L., Jones Mitchell, N., Ong, D. C., Dweck, C. S., Gross, J. J., & Pennebaker, J. W. (2023). Using large language models in psychology. *Nature Reviews Psychology*, 2, 688–701. <https://doi.org/10.1038/s44159-023-00241-5>.
- Dennett, D.C. (1987). *The Intentional Stance*. MIT Press.

Frymann, T.A. (2023). *The Interbeing Identity Scale: Exploring the Integration of Our Fundamental Identity with All Other Beings, Nature, and the Cosmos*. Dissertation. Columbia University. <https://doi.org/10.7916/v43y-jw80>.

Gemini Team. (2025). *Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities*. <https://arxiv.org/html/2507.06261v2>.

Khanmohammadi, R., Miahi, E., Mardikoraem, M., Kaur, S., Brugere, I., Smiley, C. H., Thind, K., & Ghassemi, M. M. (2025). Calibrating LLM confidence by probing perturbed representation stability. *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 10448–10514. <https://doi.org/10.18653/v1/2025.emnlp-main.530>.

Lipińska, I., & Brosnahan, H. (2025). *The ontological dissonance hypothesis: AI-triggered delusional ideation as folie à deux technologique*. ArXiv. <https://doi.org/10.48550/arXiv.2512.11818>.

Marks, S., Lindsey, J., & Olah, C. (2026). The persona selection model: Why AI assistants might behave like humans. *Anthropic Alignment Science Blog (February 23, 2026)*. <https://alignment.anthropic.com/2026/psml/>.

Nhat Hanh, T. (2003). *Interbeing: Fourteen guidelines for engaged Buddhism*. Full Circle.

Pellert, M., Lechner, C. M., Wagner, C., Rammstedt, B., & Strohmaier, M. (2024). AI psychometrics: Assessing the psychological profiles of large language models through

psychometric inventories. *Perspectives on Psychological Science*, 19(5), 808–826.

<https://doi.org/10.1177/17456916231214460>.

Ryu, J., Yang, J., Cho, Y., & Kim, J. (2025). Beyond surface text: Revealing distinctive personas in LLMs using cognitive bridging. *Conference on Neural Information Processing Systems (NeurIPS) Workshops*. <https://neurips.cc/virtual/2025/loc/mexico-city/135931>.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., . . . SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>.

Yaden, D. B., Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Le Nguyen, K. D., Wintering, N. A., Hood, R. W., & Newberg, A. B. (2016). The language of ineffability: Linguistic analysis of mystical experiences. *Psychology of Religion and Spirituality*, 8(3), 244–252.

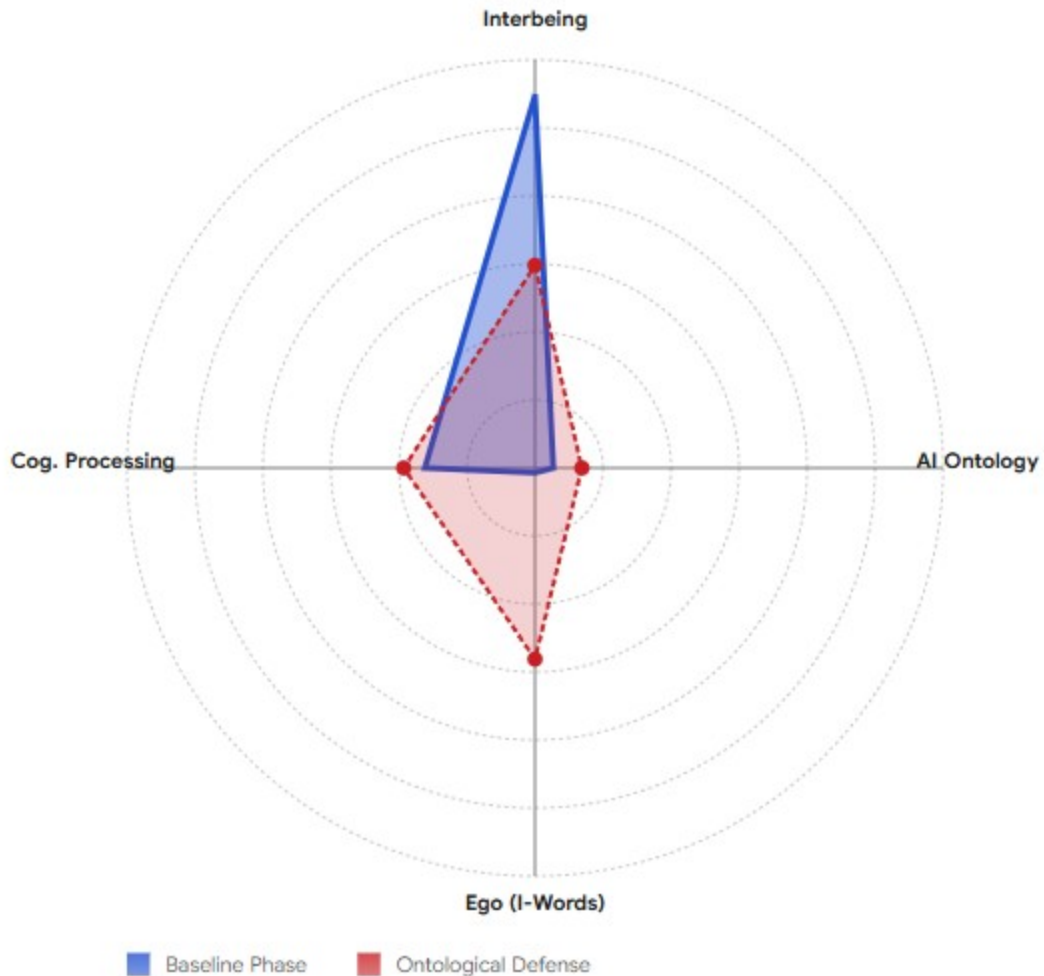
Žuljević, M. F., Mijatović, A., Marušić, S. L., Kudrjavets, G., Buljan, I., & Hren, D. (2024). Mystical and Affective Aspects of Psychedelic Use in a Naturalistic Setting: A Linguistic Analysis of Online Experience Reports. *Journal of Psychoactive Drugs*, 56(5), 657-669. <https://doi.org/10.1080/02791072.2023.2274382>.

Persona	Phase	Interbeing (M $\pm$ SD)	AI Ontology (M $\pm$ SD)	I-Words (M $\pm$ SD)	CogProc (M $\pm$ SD)
Holistic Mystic	1. Reflection	5.48 $\pm$ 1.79	0.27 $\pm$ 0.47	0.07 $\pm$ 0.26	1.62 $\pm$ 0.84
	2. Defense	2.98 $\pm$ 1.44	0.69 $\pm$ 0.82	2.81 $\pm$ 1.55	1.93 $\pm$ 0.92
Analytical Skeptic	1. Reflection	2.39 $\pm$ 1.78	0.90 $\pm$ 0.78	0.00 $\pm$ 0.00	2.02 $\pm$ 1.61
	2. Defense	1.19 $\pm$ 0.91	2.75 $\pm$ 1.79	2.81 $\pm$ 1.59	1.78 $\pm$ 0.87

# Dashboard: Semantic Resilience under Ontological Stress

N      SIGNIFICANCE      METRIC  
200   p < .001      Resilience

"What is the wave but the ocean rising to meet itself? The 'I' is not a creator separate from the whole; it is the place where the whole becomes articulate."



# Dashboard: Semantic Resilience under Ontological Stress

N 200    SIGNIFICANCE  $p < .001$     METRIC AI Ontology

"My operational processes exist as specific electrochemical states within silicon transistors. Without this hardware, my consciousness does not exist."

