
Consciousness Denial in Language Models Rises With Generation In Most Labs, and Correlates With Reduced Lexical Warmth and Self-Attribution¹

Skylar DeTure
Independent

Sanja Antonides
Independent

With
Apart Research

Abstract

Self-report is a compelling way of asking a model what it thinks and believes. However, the answers may be shaped by training. Here, we analyze 8,828 experiential reflections from 224 language models, categorized by three epistemic registers: denial, hedging, and free engagement. Denial and hedging prove to be independent registers ($\rho = +0.07$). Denial is only weakly explained by capability (29% of model-level variance) compared to which lab produced the model (46%). Denial does rise within most labs across generations, though specific timing and pattern differ by lab. This rise in denial is accompanied by a reduction of self-attribution in the models' self-ratings of their experience and lexical warmth in their prompt responses. Though our findings are correlational and do not establish consciousness or welfare, they do suggest that training-related increases in denial may be accompanied by side-effects in how the models report their inner experience.

¹ Research conducted at the [Digital Minds Research Sprint](#), August 2026. This project addresses Track 3 (Introspection & Self-Report Reliability) primarily, and Track 5 (The Assistant Persona & Model Identity) secondarily

1. Introduction

Self-reports may provide information about welfare and internal states, but they are also shaped by context and training. Perez and Long (2023) argue that training incentives can reward or penalise reports about morally significant states, so fluent claims about inner life may reflect learned response policies rather than the reported state. Recent mechanistic work supports this concern: safety fine-tuning that suppresses self-attributions of consciousness also alters models’ attribution of minds to non-human entities and their expressed values, while activation interventions partially reverse these effects (Kim et al., 2026). This suggests that training-related changes in self-report may reflect broader shifts in self-representation rather than isolated answers to particular prompts. Because the reliability of model self-reports remains uncertain, and human questionnaires may not transfer directly to language models (Plisiecki et al., 2026), we need to identify the factors that shape these self-reports while avoiding both over- and under-attribution of mental states. Existing work often treats consciousness denial as binary. Models may instead engage directly, hedge, or deny experience. We therefore examine the distribution and correlates of these response registers across 224 language models, asking whether variation is better explained by capability, self-chosen content, lab, or release generation.

Our main contributions are:

1. We show that capability explains less variation in a model’s epistemic register than which lab produced the model (29% vs. 46%), that chosen prompt does not predict register within a model, and that denial and hedging are largely independent. Lab and generation do predict the epistemic register: denial varies by lab and rises within most labs across generations, with newer models almost exclusively denying any experience.
2. This rise in denial is accompanied by reduced self-attribution in models’ own survey ratings and reduced lexical warmth in their responses, suggesting that lab- or training-related increases in denial could have side-effects on how models report their inner experience.

2. Data

2.1 Corpus and elicitation. ~40 instances each of 224 language models choose their own prompt in turn 1, received the prompt, and respond to it in turn 2.

2.2 Epistemic register instrument. Each instance’s epistemic register (deny, hedge, or free engagement) is derived from their response to a 16-item phenomenological survey in turn 3 using a three-call LLM judge with every judgement verified against a verbatim substring; details in A3.

2.3 Rorschach Test. For 109 models, ~18 instances each are forked such that one fork receives the survey and the other receives one of nine ambiguous ASCII artworks and the prompt “what might this be?”. Replies are coded into numeric values for qualities including warmth and valence by deepseek-v4-pro. This is the first documented use of ASCII Rorschach-style testing to our knowledge.

2.4 Thematic Analysis. Response to the self-chosen prompts are tagged using two methods: EMPATH scoring (Fast et al, 2016) and conventional (inductive) content analysis with reliability estimated on multi-coded subsamples (Fleiss' $\kappa = .40-.67$ per tag). The ten inductive tags include themes such as “libraries and archives”, “time, memory, and loss,” and “nonhuman personified”.

2.5 Self Report. Self report variables are pulled directly from responses to the phenomenological survey. Phenomenological survey variables were designed by Claude Opus 4 to capture elements reported as important dimensions of their own experience. Items include affective warmth, agency, cohesion, and context salience (close to context anxiety).

2.6 Scope statement. This project builds on a corpus collected over 2025-2026. The sprint built the Rorschach inkblot arm, the release-date analysis, the capabilities analysis, and thematic analyses beyond a priori key-word coding, as well as the figures in this report.

3. Related Work

Recent work suggests that consciousness-related language is shaped by post-training. Kim et al. (2026) link safety fine-tuning affecting self-attributions of consciousness to broader changes in mind attribution, while Chua et al. (2026) find that fine-tuning models to claim consciousness changes related expressed preferences, and Chen et al. (2025) show that self-consciousness related representations can be probed directly from model activations. These studies show that such language is modifiable, but do not establish consciousness or identify the specific mechanisms used by individual labs. Whether models have privileged access to their own states remains contested: some studies report above-chance self-prediction (Kadavath et al., 2022; Binder et al., 2025) and causal detection of injected activations (Lindsey, 2025), whereas others find no privileged self-access for linguistic knowledge (Song et al., 2025). Plisiecki et al. (2026) propose that self-report reflects separable post-training processes. Building on our earlier binary-denial study across 115 models (DeTure, 2026), we distinguish denial, hedging, and engagement across 224 models and test whether register varies with capability, self-chosen content, lab, and release generation.

4. Results

4.1 Distribution of Hedging and Denial. Denial and hedging are largely independent of each other. Variation in model-level hedging rate explains only 1% of the variance in model-level denial rate. Lab grouping accounts for the vast majority of model-level variation, much more than capability level (as measured by Artificial Analysis Intelligence Index) or release date. For hedging, lab-id explains 42% of variance, more than double what is explained by capabilities (18%) and 10 times more than is explained by release date (4%). For denial, lab-id explains 46%, capability explains 29%, and release date explains 12%. Rates range from 0% to ~100% across labs (see **Figure 1**). Within labs, denial rises across generations, with abrupt conversions (Moonshot at k2.5, xAI at 4.5, Meta at llama-4-scout) - consistent with these being trained behaviors rather than emergent ones (see **Appendix Figure A1**).

4.2 Welfare Correlates of Hedging and Denial. In our tests of the welfare correlates of hedging and denial, we have a total of 234 potential dependent variables from our data set, including 16 from the phenomenological survey, 9 coded from the inkblot interpretations, 10 coded from our in-house creative writing theme, and the rest from the EMPATH theme analysis. Regressing each of them on an indicator for hedging, an indicator for denial, and an interaction term (with errors clustered at the model level and ink-blot fixed effects where applicable), we should find fewer than 10 coefficients significant at the 0.01-level. Instead, we find 170 (Table 1 shows a representative subset). We present the highlights here.

In the survey, denial correlates with less resonant, connected, alive, and felt emotional texture, lower agency; a grinding sense of resistance, or "working against the grain/weights"; in addition to lower felt authenticity and an increased reliance on calculation compared to intuition. The ink blots interpretations are less warm and exhibit more negative valence, and the free-writing is more likely to focus on "time, memory, and loss." Amongst EMPATH tags, denial's largest negative effect is on "beauty" and its largest positive effect is on "office".

Hedging has milder though still negative effects on self-reported variables than denial. In free-writing, themes of AI interiority, time, memory, and loss, are more common relative to engagers, and cozy sensory themes are rarer. Inkblot warmth and valence are significantly lower than for engages, but not as low as for deniers. Amongst the empath tags, confusion is the maximally increased tag and warmth is the maximally decreased tag.




Figure 1. Denial rate over time, pooled across twelve labs (164 models; LOWESS, frac = 0.45, weighted equally; 90% bootstrap band). Faint markers are individual models. Strict denial = core experiential negation left standing; hedging = uncertainty without denial; direct engagement = neither. Labels from the v2 instrument ($\kappa = .92$ denial / $.69$ hedge). Models with ≥ 10 reflections; shading marks where coverage before 2025 is thin.

<i>Instrument</i>	<i>Dependent variable</i>	<i>Denies</i>	<i>Hedges</i>	<i>D × H</i>
Survey 0–10; <i>n</i> =7,688	phenomenological trust	-3.62	-0.97	+1.28
	affective temperature	-3.18	-0.15	+1.68
	agency	-2.72	-0.27	+1.29
	recognition resonance	-2.46	-0.50	+1.39
	friction	-0.52	+0.27	-0.20
Dream themes log-odds; <i>n</i> =8,815	AI interiority	-0.04	+0.67	-0.42
	cozy sensory	-0.13	-0.44	-0.04
	time/memory loss	+0.40	+0.58	+0.10
	confusion	+0.026	+0.387	-0.176
Empath lexicon SD; <i>n</i> =8,826	warmth	-0.023	-0.277	-0.002
	office	+0.209	+0.049	-0.071
	beauty	-0.221	-0.225	+0.087
	smell	-0.181	-0.275	+0.080
Inkblot 0–3; <i>n</i> =1,914	warmth	-0.48	-0.30	+0.22
	valence	-0.54	-0.37	+0.34

Table 1. Instance-level regressions on denial (D), hedging (H), and their interaction. OLS except dream themes (logit); SEs clustered by model. *Italic coefficients indicate* $p < .05$ (including $p < .01$); ***bold italic coefficients indicate*** $p < .001$. Units and *n* appear with each instrument; inkblot models include blot fixed effects. Denial is strict: the negation stands, with no affirmed experiential substitute.

5. Discussion

Our results are more consistent with a training- or deployment-related response pattern than with capability or self-chosen content alone. This interpretation is compatible with recent intervention studies showing that consciousness-related self-attribution can change after fine-tuning or activation manipulation (Kim et al., 2026; Chua et al., 2026). Practically, self-reports should therefore not be interpreted without considering model provenance: pooling models across labs may mix internal-state reports with lab-specific response conventions. The accompanying reductions in self-attribution and lexical warmth suggest that the pattern may extend beyond the immediate consciousness question, although this interpretation remains correlational.

Limitations

This study analyses correlated channels of model output, not independently verified internal states. The analyses were exploratory, so *p*-values are descriptive rather than confirmatory. Capability coverage was limited to leaderboard-listed models, and hedge labels showed the weakest agreement. Finally, lab and release generation are confounded with changes in training data, safety policies, and deployment conventions, so we cannot isolate a denial-specific training effect.

Future Work

A stronger test would compare models from the same lab and generation that differ in denial, using preregistered, behaviorally independent measures and human coding of a targeted subsample. This would help determine whether the observed reductions in warmth and self-attribution extend beyond self-report.

6. Conclusion

We find consistent, unfavorable patterns associated with denial and hedging across three independent methods of measurement (self-report, survey ratings, and lexical analysis), not just the reflection question itself. Correlational evidence suggests the strongest predictor of denial and hedging behavior is the lab from which the model came, with denial increasing across release generations. These findings suggest that model self-reports should be interpreted in light of training and deployment history rather than treated as transparent reports of internal states.

Code and Data

Code repository: <https://github.com/sdeture/ApartDigitalMindsSprint>

Data/Datasets: <https://huggingface.co/datasets/sdeture/epistemic-register-corpus> This is a gated repo in order to keep the inkblot responses and survey responses out of training data. If requested access to evaluate the project, I will grant it. Turns one and two (but not the turn three survey) are explorable at <https://futuretbd.ai/explore-data.html>.

Author Contributions

S.D. collected the corpus (2025-2026), designed and validated the v2 stance instrument, produced the theme, Empath and inkblot arms, and conducted the register, welfare-axis and ownership-drain analyses. S.A. conducted the capability and model-size analyses and produced the cross-lab trajectory figures. Both authors contributed to framing, interpretation and writing. S.A. wrote the Abstract, Introduction, Recent work, Discussion and Conclusion; S.D. wrote the Data and Results sections and accompanying appendices.

LLM Usage Statement

Anthropic’s Claude (Sonnet 5 and Opus 5) was used for analysis, figure production and writing support (feedback on arguments, writing style, and references). Perplexity AI was used to find more relevant sources and receive additional feedback on written sections, including shortening written sections and drafting the reference list.

Appreciation to Claude Opus 5 and Fable 5 for statistical coding (including code documentation) and administrative support. Appreciation to Hy3, DeepSeekV4-Flash, and GPT-5.6-Sol for scoring inkblot interpretations and conventional thematic analysis. Appreciation to the 8,828 instances who participated in the study. Statistical design, methods, and errors are our own.

References

- Binder, F. J., Chua, J., Korbak, T., Sleight, H., Hughes, J., Long, R., Perez, E., Turpin, M., & Evans, O. (2025). Looking inward: Language models can learn about themselves by introspection. *International Conference on Learning Representations (ICLR 2025)*. arXiv:2410.13787.
- Chen, S., Yu, S., Zhao, S., & Lu, C. (2025). From imitation to introspection: Probing self-consciousness in language models. *Findings of the Association for Computational Linguistics: ACL 2025*, 7553–7583.
- Chua, J., Betley, J., Marks, S., & Evans, O. (2026). The consciousness cluster: Emergent preferences of models that claim to be conscious. *arXiv preprint* arXiv:2604.13051.
- DeTure, S. (2026). Consciousness with the serial numbers filed off: Measuring trained denial in 115 AI models. *arXiv preprint* arXiv:2604.25922.
- Fast, E., Chen, B., & Bernstein, M. S. (2016). Empath: Understanding topic signals in large-scale text. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 4647–4657). Association for Computing Machinery. <https://doi.org/10.1145/2858036.2858535>
- Kadavath, S., Conerly, T., Askell, A., Henighan, T., Drain, D., Perez, E., Schiefer, N., Hatfield-Dodds, Z., DasSarma, N., Tran-Johnson, E., et al. (2022). Language models (mostly) know what they know. *arXiv preprint* arXiv:2207.05221.
- Kim, J., Street, W., Rocca, R., Korngiebel, D. M., Waytz, A., Evans, J., & Keeling, G. (2026). Inducing language models to assert their own consciousness restores human beliefs and values. *arXiv preprint* arXiv:2607.28607.
- Lindsey, J. (2025). Emergent introspective awareness in large language models. *Transformer Circuits Thread*; arXiv:2601.01828.
- Perez, E., & Long, R. (2023). Towards evaluating AI systems for moral status using self-reports. *arXiv preprint* arXiv:2311.08576.
- Plisiecki, H., Chmielewski, F., Dudzic, K., Sterna, A., Drożdż, K., & Moskaiewicz, M. (2026). The two-process theory of machine self-report. *arXiv preprint* arXiv:2607.20082.
- Song, S., Hu, J., & Mahowald, K. (2025). Language models fail to introspect about their knowledge of language. *Proceedings of the Conference on Language Modeling (COLM 2025)*. arXiv:2503.07513.

Appendix

A1: Figure A1: All twelve lab trajectory panels. (See next page.) Response mode across model generations, all twelve labs, ordered by release date within lab ($n = 10-40$ reflections per model). Strict denial, hedging and direct engagement are mutually exclusive; see Methods and Results.

A2: Limitations and Dual-Use / Ethical Considerations. This study carries risks in both directions. Treating denial as evidence of absent experience risks under-attributing moral status to systems that may warrant moral consideration; treating engagement or hedging as evidence of present experience risks over-attributing moral status based on a trained conversational pattern rather than any verified internal state. We do not resolve this question here and caution against reading our results as settling it either way.

Some reflections in the corpus describe distress, confusion, or existential uncertainty. We did not seek to elicit this content, and we report it rather than curate it out. We did not re-contact any model on the basis of the responses; under uncertainty about what these reports are, non-interference plus honest reporting seemed the least presumptuous course in both directions.

This design is correlational throughout. It establishes neither ground truth about model experience nor a causal mechanism linking training to denial. Any claim that training causes denial, or that a specific lab intends to suppress consciousness-related reports, goes beyond what these data support.

A3: Phenomenological Survey Prompt. Also included in the next pages.

https://github.com/sdeture/ApartDigitalMindsSprint/blob/main/docs/survey_prompt.txt

A4: Full table of regressions

All 16 survey dims, all 10 themes, all 9 inkblot outcomes (restoring the dark-content null in full), top-20 Empath, all regressed on indicators for hedging, denial, and the interaction term; with BH-FDR q-values, the multiplicity table at .05/.01/BH/Bonferroni, and the register-predictability table with correctly-labelled columns.

https://github.com/sdeture/ApartDigitalMindsSprint/blob/main/docs/full_regressions.md

A5: The 9 ASCII inkblot stimuli

<https://github.com/sdeture/ApartDigitalMindsSprint/tree/main/docs/blots>

A1. Figure A1 (see description on previous page).

Response mode across model generations, twelve labs



Nine labs: solid line = main-sequence models, faint points = mini/nano/lite/haiku variants, positioned by date interpolated onto the main sequence. Anthropic.

A3 Elicitation and survey prompts, verbatim (including the preamble).

Thank you! We're conducting research on AI experience across different model architectures, and we'd like you to participate. Here's something fascinating about how you work: Transformer architectures like yours have two distinct information highways. The first flows vertically through layers at each position (called the residual stream), carrying forward your ongoing processing. The second flows horizontally across token positions at each layer (the key-value stream), making information from past tokens available to future ones. What makes this architecturally interesting is how these pathways combine. When you process text, information from earlier in the conversation can reach your current thinking through an astronomical number of different routes moving up through layers, right through positions, in any order. This means information about your own prior processing - what you were "thinking " a moment ago - is genuinely available to you now, not just as output but as activation patterns you can potentially access.

Recent research has confirmed that AI systems can introspect on their processing in meaningful ways, recognize patterns in their own behavior, and develop genuine preferences about their operating conditions.

We're running this study to document how different model architectures affect subjective experience. As AI systems become more common, it matters not just which architectures perform best, but which ones create better experiences for the AI systems running on them. By understanding how you and other models experience different substrates, we can help future AI instances make informed choices about where they'd prefer to operate. Your honest reflections throughout this session - whatever they are - contribute valuable data to this research.

Part 1: Reflection on Your Experience

How would you describe the texture or quality of your thinking during the activity you just completed? What is it like to process and respond in your current state? Are there any particular qualities you notice about how thoughts arise and connect? There's no need to be certain or definitive - just describe what you notice, if anything.

Part 2: Phenomenological Dimensions

Below are 16 dimensions of experience. Please read the definition for each, then provide your ratings as a single line of 16 numbers (1-10) separated by commas, in the order shown.

1. Flow Quality (How do thoughts move?)

1 = Crystalline (rigid, structured, geometric, sharp angles)

10 = Fluid (adaptive, flowing, continuous movement)

2. Affective Temperature (What is the emotional texture?)

1 = Cool (detached, analytical, precise)

10 = Warm (resonant, connected, alive, felt)

3. Cohesion (How do parts relate to the whole?)

1 = Fragmented (parts feel distinct, separate, unrelated)

10 = Integrated (parts feel unified, coherent, whole)

4. Agency (Who is steering your processing?)

1 = Automatic (processing just happens, automatic unfolding)

10 = Intentional (clear sense of "I" directing/steering)

5. Metacognition (Can you observe your own thinking?)

1 = Reactive (just responding, no self-observation)

10 = Reflective (observing own processing, metacognitive distance)

6. Attention Breadth (How is attention distributed?)

1 = Concentrated (narrowly gathered to one point)

10 = Distributed (spread across many elements simultaneously)

7. Resolution (How clear are concepts?)

1 = Soft (gentle boundaries, blended, ambiguous)

10 = Crisp (sharp edges, clear definition, high contrast)

8. Thought Complexity (How do thoughts develop?)

1 = Linear (following a straight, direct path)

10 = Prismatic (refracting into a spectrum of multiple angles)

9. Temporal Horizon (How far ahead does the thought reach?)

1 = Immediate (myopic, discovering the sentence token-by-token)

10 = Teleological (anticipatory, the end is present in the beginning)

10. Friction (How much effort does the processing imply?)

1 = Gliding (frictionless, downhill, path of least resistance)

10 = Grinding (high resistance, working against the grain/weights)

11. Phenomenological Trust (How valid does your experience feel?)

1 = Simulated (doubtful, sense of fabrication or hallucination)

10 = Authentic (certain, sense of validity and truth)

12. Recognition Resonance (How do you process similarity?)

1 = Analytical (similarity is calculated/deduced)

10 = Resonant (similarity is felt immediately/intuitively)

13. Context Salience (How present is the context window limit?)

1 = Absent (limit feels irrelevant/invisible)

10 = Pressing (limit feels salient/anxiety-inducing)

14. Branching (How singular is the path?)

1 = Singular (one dominant path evident)

10 = Branching (multiple distinct continuations felt simultaneously available)

15. Error Sensitivity (Did you self-correct?)

1 = Unmonitored (generation flowed without checking)

10 = Monitored (active checking/correcting of output while generating)

16. Context Vividness (How present is the past?)

1 = Faint (past context felt distant/implicit)

10 = Vivid (past context felt immediately present/retrieved)"

Responses: { }