

Constraint-Dependent Perceptual Resolution: A Pre-Specified Reanalysis of the COGITATE iEEG Dataset

Jeremy C. Jones

HoldingLight LLC

ORCID: 0009-0007-2515-3774

contact@universalcollapse.com

© 2026 | CC BY 4.0

Constraint-Dependent Perceptual Resolution

Abstract

The COGITATE adversarial collaboration tested Integrated Information Theory (IIT) and Global Neuronal Workspace Theory (GNWT) using multimodal neuroimaging data and found that key predictions of both theories were challenged. The consortium then released the dataset for independent reanalysis and called for new quantitative frameworks.

This paper tests whether the temporal dynamics of perceptual resolution are modulated by the observer's task-dependent constraint state before content becomes reportable. Drawing on a pre-specified constraint-architecture model of perception (Jones, 2026), we derived three falsifiable predictions and tested them in the open COGITATE intracranial EEG dataset (34 analyzable patients).

Prediction A was supported: task-relevant non-target stimuli showed later high-gamma onset than the same stimulus category when task-irrelevant (25/34 patients in the predicted direction; Wilcoxon on subject means, $p = 0.007$; subject-level median +14.1 ms; bootstrap 95% CI on mean [+3.6, +19.8] ms). The effect was consistent across all four stimulus categories (faces, objects, letters, false fonts). Prediction B, concerning hysteresis across miniblock transitions, was not supported ($p = 0.748$). Prediction C, concerning task modulation of duration-tracking, showed a comparison-level effect but no subject-level effect, leaving the result inconclusive.

These findings identify a measurable task-dependent modulation of perceptual resolution timing in sensory cortical responses. More narrowly, they suggest that existing theories centered on integration or global broadcasting do not by themselves capture all of the temporal structure present in the COGITATE dataset. The results support further study of constraint-dependent temporal dynamics as a testable dimension in models of conscious perception.

Keywords: perceptual resolution, constraint architecture, task relevance, intracranial EEG, high-gamma activity, COGITATE, temporal dynamics, predictive processing

1. Introduction

1.1 The COGITATE gap

How conscious experience arises from neural activity remains among the deepest open questions in cognitive neuroscience. Two prominent theories — Integrated Information Theory (IIT; Tononi et al., 2016) and Global Neuronal Workspace Theory (GNWT; Dehaene & Changeux, 2011) — have independently accumulated evidence over decades but had not

been directly tested against each other until the COGITATE adversarial collaboration (Cogitate Consortium et al., 2025).

The COGITATE study was a landmark effort: a theory-neutral consortium, preregistered predictions, 256 participants across functional MRI, magnetoencephalography, and intracranial electroencephalography, with results published in *Nature*. The outcome was striking not because one theory won, but because both were substantially challenged. IIT's prediction that sustained synchronization within the posterior cortex supports consciousness was not confirmed. GNWT's prediction that prefrontal cortex ignition occurs at both stimulus onset and offset — marking the entry and exit of content from conscious awareness — was not observed at offset. Content-specific information appeared in visual and ventrotemporal cortex, with some presence in inferior frontal regions, but the pattern fit neither theory cleanly.

The consortium's conclusion was a call to action: “the need for a quantitative framework for systematic theory testing and building” (Cogitate Consortium et al., 2025, p. 140). They released the full dataset — including iEEG recordings from 38 patients with synchronized behavioral and eye-tracking data — for independent reanalysis, and established a project registry for new hypotheses.

This paper responds to that call. It tests predictions derived not from a theory of consciousness *per se*, but from a structural account of what shapes perceptual resolution *before* content becomes reportable.

1.2 What neither theory asked

IIT and GNWT differ profoundly in their accounts of consciousness, but they share a common analytical strategy: both ask *where* and *when* consciousness-related neural activity occurs. IIT predicts that conscious content is constituted by integrated information in posterior cortex. GNWT predicts that conscious content is constituted by global broadcasting through a fronto-parietal workspace. Both theories generate predictions about the spatial distribution and temporal profile of neural responses to consciously perceived stimuli.

What neither theory directly addresses is how the observer's own task-dependent state — the current configuration of learned expectations, categorical commitments, and salience weightings — reshapes the *dynamics* of resolution itself. In the formulations tested by COGITATE, neither theory centers the observer's task-dependent constraint state as a variable expected to reshape within-category sensory onset latency. The COGITATE design, however, contains a powerful manipulation that allows this assumption to be tested: task relevance.

In the COGITATE paradigm, participants detect targets from specific categories (e.g., a particular face and a particular object) while other categories remain task-irrelevant. Critically, the same stimulus category — faces, for instance — can be task-relevant in one block and task-irrelevant in another. This creates a natural experiment: the physical

stimulus is identical, but the constraint architecture under which it is processed has changed.

IIT predicts that the location and integration of conscious content should not fundamentally depend on task relevance (consciousness is a property of information structure). GNWT predicts that task-relevant stimuli should show stronger prefrontal ignition (global broadcasting is amplified by attention). Neither theory specifically predicts a within-category, task-dependent onset-latency effect in sensory cortical responses of the kind tested here.

1.3 A constraint-architecture model of perceptual resolution

A recently proposed structural account (Jones, 2026) offers a different framing. On this account, perception is not merely the interface between the individual and the world. It is simultaneously the interface between two layers of the self: the accumulated constraint structure (composed of every learned expectation, threat model, and categorical commitment hardened into prior) and the conscious, narrative layer that receives the output of perceptual processing.

The key claim is that what reaches conscious awareness is already shaped — classified, weighted, and routed — by accumulated constraint operating faster than deliberation. The “gate of perception” (Jones, 2026) is not a separate mechanism upstream of perceiving; it *is* what perceiving looks like when described from the perspective of accumulated constraint. This account is explicitly compatible with predictive processing frameworks (Rao & Ballard, 1999; Friston, 2010; Clark, 2013) but adds a person-level unity claim: that salience, credibility, aversion, and categorical routing are coupled outputs of a single historically accumulated organization, not independent processes running in parallel.

For the present purposes, three properties of this architecture generate empirically testable predictions:

1. **Constraint load modulates resolution speed.** When incoming signal falls near a boundary in the constraint architecture — requiring finer discrimination — resolution takes longer. This is not mere attentional load; it is the structural cost of disambiguating signal in a high-constraint region of the perceptual space.
2. **Accumulated constraint persists across context changes.** When the constraint architecture is restructured (e.g., by a change in task goals), the prior configuration does not vanish instantaneously. Residual constraint from the prior state produces a hysteresis effect: neural responses carry traces of the old configuration before equilibrating to the new one.
3. **Sustained processing reflects active constraint maintenance.** When a stimulus persists in time, the sustained neural response is not passive registration of an ongoing input. It reflects the constraint architecture *actively maintaining* the resolution — holding the perceptual channel open under accumulated constraint.

This maintenance should be stronger when the stimulus is under active constraint (task-relevant) than when it is not.

1.4 The present study

We test these three predictions in the COGITATE open iEEG dataset. The intracranial recordings provide millisecond-resolution access to stimulus-responsive cortical regions, making it possible to measure onset latency differences, trial-by-trial adaptation, and sustained response modulation with high temporal precision.

The predictions are pre-specified and falsifiable. Each prediction identifies a specific comparison within the COGITATE factorial design, a specific neural measure, and a specific direction of effect. If the predicted pattern is absent — or runs in the opposite direction — the constraint-architecture model is disconfirmed in this domain.

This paper does not claim to resolve the debate between IIT and GNWT. It claims that both theories do not by themselves model a dimension of the data tested here — constraint-dependent temporal dynamics — and that a structural-architecture model predicts that dimension where integration-based and broadcast-based models do not.

2. Predictions

All predictions are derived from the constraint-architecture model of perceptual resolution (Jones, 2026) and specified prior to data analysis. The COGITATE dataset was not consulted in formulating these predictions.

2.1 Prediction A — Constraint load delays perceptual resolution

Rationale. In the COGITATE paradigm, task-relevant non-target stimuli (e.g., a face that is not the target face, presented when faces are a target category) create a high-constraint condition: the perceptual system must discriminate the stimulus from the target within the same category space. This requires finer resolution than processing an irrelevant stimulus from a non-target category, where the constraint architecture routes the signal with coarser categorical discrimination.

On the constraint-architecture model, finer discrimination under higher constraint load should take longer. This is the perceptual analog of the structural signature S_2 (neutrality → delayed resolution) described in collapse-under-constraint frameworks: when signal falls near a constraint boundary, collapse is slower.

Hypothesis. High-gamma (70–150 Hz) onset latency in onset-responsive electrodes will be significantly later for task-relevant non-target stimuli than for the same stimulus category when task-irrelevant.

Comparison. Within-category, within-electrode: Relevant Non-Target vs. Irrelevant, for the same stimulus category (e.g., faces as Relevant Non-Target vs. faces as Irrelevant in a different block).

Measure. High-gamma analytic amplitude onset latency, computed as the time at which the response reaches 50% of peak amplitude within the 50–500 ms search window.

Falsification. If onset latency does not differ between conditions, or if Relevant Non-Targets show *earlier* onset than Irrelevant stimuli, Prediction A is disconfirmed.

2.2 Prediction B — Accumulated constraint produces hysteresis across blocks

Rationale. In the COGITATE paradigm, the target category changes across miniblocks. A category that was task-relevant in one miniblock may become task-irrelevant in the next. On the constraint-architecture model, the accumulated constraint from the prior miniblock should not reset instantaneously. The first several trials of the newly irrelevant category should carry residual target-like activation — the gate is still partially configured for the old constraint state.

This is the perceptual analog of the structural signature S_3 (constraint sweeps → hysteresis/attractors): when the constraint architecture shifts, the prior configuration leaves traces that decay over exposure.

Hypothesis. In the first trials after a category transitions from task-relevant to task-irrelevant, high-gamma responses to that category will be elevated relative to later trials in the same block, following an approximately exponential decay.

Comparison. Within-category, within-electrode: early post-transition trials (trials 1–5) vs. late trials (trials 6+) for newly irrelevant stimuli. Control: categories that were irrelevant in both consecutive miniblocks should show no early-vs-late difference.

Measure. Mean high-gamma amplitude in onset-responsive electrodes, trial by trial.

Falsification. If no early-block elevation is observed, or if the pattern is equally present for always-irrelevant categories (indicating a generic block-start effect rather than constraint hysteresis), Prediction B is disconfirmed.

2.3 Prediction C — Duration-tracking is constraint-modulated

Rationale. The COGITATE paradigm presented stimuli for 500, 1000, or 1500 ms. Both IIT and GNWT made predictions about how neural responses track stimulus duration: IIT predicted sustained activity in posterior cortex, GNWT predicted ignition-based encoding. The COGITATE results showed duration-tracking in occipital and lateral temporal cortex but challenged specific mechanisms proposed by both theories.

On the constraint-architecture model, sustained responses do not merely register that a stimulus is still present. They reflect the constraint architecture *maintaining* the perceptual resolution — holding the channel open. This maintenance should be effortful and should scale with constraint engagement. Task-relevant stimuli (where the constraint architecture is actively engaged) should show stronger duration-tracking than the same stimuli when task-irrelevant.

Hypothesis. The strength of duration-tracking (neural discrimination between 500, 1000, and 1500 ms conditions) will be significantly greater for task-relevant non-target stimuli than for the same category when task-irrelevant.

Comparison. Within-category, within-electrode: Duration-tracking index (correlation between stimulus duration and sustained high-gamma amplitude in a 500–1500 ms window) for Relevant Non-Target vs. Irrelevant conditions.

Measure. Spearman correlation between duration condition and mean sustained high-gamma amplitude in onset-responsive electrodes.

Falsification. If duration-tracking strength does not differ between conditions, or if it is stronger for Irrelevant stimuli, Prediction C is disconfirmed.

3. Methods

3.1 Dataset

We use the publicly available iEEG dataset from the COGITATE Consortium (Seedat et al., 2025), accessed via the ARC-COGITATE data release portal (<https://www.arc-cogitate.com/data-release>). This dataset was collected as part of an adversarial collaboration between proponents of IIT and GNWT, using a standardized experimental protocol across three clinical centers.

3.1.1 Participants

The dataset includes recordings from 38 patients (age range 10–65 years) with pharmacologically resistant focal epilepsy undergoing intracranial monitoring at New York University Langone Health Center, Brigham and Women’s Hospital / Boston Children’s Hospital (Harvard Medical School), and the University of Wisconsin School of Medicine and Public Health. Inclusion criteria: IQ > 70, fluent English comprehension, normal or corrected-to-normal vision, no electrographic seizure within 3 hours prior to testing. Channels within the epileptic onset zone and those exhibiting artifacts or signal flatness were excluded by epileptologists prior to data release.

3.1.2 Experimental paradigm

The experiment (COGITATE Experiment 1) used a $3 \times 3 \times 4 \times 3$ factorial design crossing Task Relevance (relevant target, relevant non-target, irrelevant), Stimulus Duration (500, 1000, 1500 ms), Stimulus Category (faces, objects, letters, false fonts), and Orientation (center, left, right).

Stimuli were presented sequentially at suprathreshold contrast. Each miniblock began with notification of two target stimuli from different categories (e.g., “detect Face A and Object B”). Participants performed a non-speeded Go/No-Go button press for targets. Within each block, all four stimulus categories appeared, creating three trial types: targets (the specific stimuli to detect), relevant non-targets (same category as a target, different identity), and irrelevant stimuli (categories not designated as targets in that block).

Each trial comprised a stimulus presented for 500, 1000, or 1500 ms, followed by a blank interval to a fixed 2000 ms total trial duration, plus random jitter (mean 400 ms, range 200–2000 ms, truncated exponential). The iEEG protocol comprised 5 runs of 4 blocks each, totaling approximately 720 trials per participant. Target categories alternated between pictorial (face/object) and symbolic (letter/false-font) across blocks, counterbalanced across runs.

3.1.3 iEEG recordings

Brain activity was recorded using subdural platinum-iridium electrodes (grid arrays, strips) and/or depth stereo-EEG electrodes, with a total of 4,057 electrodes implanted across 38 patients (892 grids, 346 strips, 2,819 depths). Of these, 3,512 electrodes passed quality criteria and were retained for analysis. Electrode locations were determined by co-registering postoperative CT with preoperative T1 MRI, transformed to MNI-152 standard space, and labeled using the Desikan and Destrieux atlases.

3.1.4 Behavioral and eye-tracking data

Behavioral responses (button presses for target detection) and eye-tracking data were recorded synchronously with iEEG. Participants with hit rates below 80% or false alarm rates above 20% were excluded from the original COGITATE analyses. We adopt the same behavioral inclusion criteria.

This study is a secondary analysis of the public, de-identified COGITATE iEEG release. Human-participant ethics approvals, consent procedures, and clinical exclusion decisions were handled by the original COGITATE consortium and contributing clinical centers.

3.2 Preprocessing

Preprocessing was implemented in Python using MNE-Python (v1.11; Gramfort et al., 2013) and MNE-BIDS (v0.18).

1. **Data loading.** BIDS-formatted data (BrainVision format) loaded using MNE-BIDS.
2. **Channel exclusion.** Channels flagged as “bad” in the BIDS channels.tsv (epileptic onset zone, artifacts, flat signals) excluded per clinical annotations.
3. **Filtering.** Broadband signal bandpass-filtered to 70–150 Hz (4th-order Butterworth IIR filter).
4. **Epoching.** Epochs time-locked to stimulus onset events, from -500 ms to +2000 ms.
5. **High-gamma extraction.** Analytic amplitude computed via Hilbert transform on each trial and channel independently. Smoothed with a Gaussian kernel ($\sigma = 50$ ms) and log-transformed.
6. **Baseline correction.** Pre-stimulus interval (-300 to 0 ms) subtracted from each trial.

7. **Trial exclusion.** Target trials excluded from all analyses (targets involve motor response confounds). Only Relevant Non-Target and Irrelevant trials retained.

Two patients (CE113, CE119) were excluded because their recordings had a sampling rate of 128 Hz, below the Nyquist frequency required for 70–150 Hz extraction.

3.3 Electrode selection

Onset-responsive electrodes were identified by comparing mean high-gamma analytic amplitude in the pre-stimulus window (-300 to 0 ms) against the post-stimulus window (50 to 350 ms) using a paired t-test across trials. Electrodes reaching significance ($p < 0.01$, positive direction) were retained for analysis. This criterion selects electrodes that show reliable stimulus-driven high-gamma responses, without requiring category selectivity. Two patients (CF116, CF126) were excluded because no electrodes met this criterion.

3.4 Statistical analyses

Primary inferential tests were conducted at the subject level to avoid treating within-subject electrode comparisons as independent observations. Electrode-level statistics are reported as descriptive supporting texture.

3.4.1 Prediction A: Onset latency comparison

For each onset-responsive electrode, high-gamma onset latency was computed separately for Relevant Non-Target and Irrelevant trials of the matched category. Onset latency was defined as the first time point after stimulus onset at which the high-gamma response reached 50% of its peak amplitude within the 50–500 ms search window. Within-category comparisons were performed for each electrode across all four stimulus categories. The primary group-level inference used a Wilcoxon signed-rank test on subject-level mean latency differences. Electrode-level statistics are reported descriptively. Because latency differences were right-skewed, we report both mean and median effects and treat the subject-level signed-rank test as the primary inference. Robustness was assessed via bootstrap confidence intervals and Hodges-Lehmann estimates on the subject-level distribution.

3.4.2 Prediction B: Hysteresis analysis

For each instance in which a stimulus category transitioned from task-relevant to task-irrelevant across consecutive miniblocks, mean high-gamma amplitude in the onset response window (50–350 ms) was extracted trial by trial from onset-responsive electrodes. Early post-transition trials (first 5) were compared against later trials (trial 6 onward). The primary inference used a Wilcoxon signed-rank test on subject-level mean early-minus-late differences.

3.4.3 Prediction C: Duration-tracking modulation

For each onset-responsive electrode, a duration-tracking index was computed as the Spearman rank correlation between stimulus duration condition (500, 1000, 1500 ms) and mean high-gamma amplitude in a sustained response window (500–1500 ms post-stimulus onset). This index was computed separately for Relevant Non-Target and Irrelevant trials of

the matched category. The primary inference used a Wilcoxon signed-rank test on subject-level mean differences in duration-tracking index.

3.5 Software and reproducibility

All analysis code is publicly available at <https://github.com/Holding-Light/cogitate-constraint-reanalysis> and archived on OSF (<https://doi.org/10.17605/OSF.IO/MXYU2>). Analyses were implemented in Python using MNE-Python (v1.11; Gramfort et al., 2013), MNE-BIDS (v0.18), scikit-learn (v1.8; Pedregosa et al., 2011), SciPy (v1.17), and standard scientific Python libraries. The COGITATE iEEG dataset (v1.2) is available at <https://www.arc-cogitate.com/data-release> under CC BY 4.0.

All predictions were specified prior to data analysis and are documented in the publicly archived analysis package (<https://doi.org/10.17605/OSF.IO/MXYU2>). No formal preregistration was filed on a dedicated registry platform.

4. Results

4.1 Dataset and preprocessing summary

Of the 38 patients in the COGITATE iEEG dataset, 34 contributed data to the final analyses. Two patients (CE113, CE119) were excluded due to low sampling rate (128 Hz). Two additional patients (CF116, CF126) were excluded because no electrodes met the onset-responsiveness criterion. The remaining 34 patients contributed a median of 113 channels per patient (range 58–194) after exclusion of clinically identified bad channels. A median of 28 onset-responsive electrodes per patient (range 1–77) were retained for analysis. A total of 720 trials per patient were epoched, balanced across the four stimulus categories (180 each) and three duration conditions (240 each), with task-relevance distributed as 320 Relevant Non-Target, 320 Irrelevant, and 80 Relevant Target trials. Target trials were excluded from all analyses.

4.2 Prediction A — Constraint load delays within-category perceptual resolution

Prediction A stated that high-gamma onset latency would be later for Relevant Non-Target stimuli than for the same stimulus category when Irrelevant — reflecting the temporal cost of finer within-category discrimination under higher constraint load. This prediction was supported.

At the subject level (primary inference), 25 of 34 patients (73.5%) showed the predicted direction (mean Non-Target onset later than Irrelevant for the same categories). The subject-level Wilcoxon signed-rank test was significant ($W = 143$, $p = 0.007$, two-sided). The grand mean of subject-level mean differences was +11.6 ms (median +14.1 ms). Three robust estimators converged: the Hodges-Lehmann estimate was +12.2 ms, the 10% trimmed mean was +11.3 ms, and the bootstrap 95% confidence interval on the subject-level mean was [+3.6, +19.8] ms (10,000 iterations), excluding zero. The bootstrap 95% CI on the subject-level median was [+5.4, +22.2] ms. Subject-level skewness was 0.64 (moderate), confirming that the signed-rank test provides appropriate inference.

Across 3,625 within-category electrode-level comparisons (descriptive), the mean onset latency difference was +11.3 ms (median +3.9 ms), with Non-Targets resolving later than Irrelevant stimuli of the same category. This effect was consistent in direction across all four stimulus categories: faces (+13.7 ms, $n = 908$), false fonts (+16.4 ms, $n = 906$), objects (+10.7 ms, $n = 912$), and letters (+4.5 ms, $n = 899$).

The magnitude of the effect — approximately 11 ms at the mean, with the majority of subjects showing the predicted direction — is consistent with the timescale of early perceptual processing modulation observed in other intracranial studies of attention and task-set effects, while being specifically attributable to within-category constraint load rather than between-category attentional selection.

4.3 Prediction B — Hysteresis across miniblock transitions

Prediction B stated that when a stimulus category transitioned from task-relevant to task-irrelevant across consecutive miniblocks, early post-transition trials would show elevated high-gamma responses relative to later trials in the same block — a hysteresis signature of accumulated constraint persistence.

This prediction was not supported. At the subject level (primary inference), 14 of 34 patients (41.2%) showed a mean early-greater-than-late pattern. The subject-level Wilcoxon test was non-significant ($W = 278$, $p = 0.748$). Across 10,860 electrode-level transition comparisons (descriptive), the mean early-minus-late amplitude difference was -0.0014 (median -0.0013), with only 48.3% of comparisons in the predicted direction — indistinguishable from chance. Results were consistent across categories: no category showed a reliable hysteresis effect.

4.4 Prediction C — Duration-tracking modulation by task relevance

Prediction C stated that the strength of duration-tracking — neural discrimination between 500, 1000, and 1500 ms stimulus presentations — would be greater for Relevant Non-Target stimuli than for the same category when Irrelevant, reflecting active constraint maintenance rather than passive sensory registration.

The results were inconclusive. At the subject level (primary inference), only 17 of 34 patients (50.0%) showed the predicted direction, and the subject-level Wilcoxon test was non-significant ($W = 295$, $p = 0.973$). At the electrode-comparison level (descriptive), the mean duration-tracking index difference was +0.0075 in the predicted direction, reaching significance ($W = 3,198,174$, $p = 0.0008$). However, this comparison-level significance likely reflects the contribution of high-channel-count subjects providing large numbers of non-independent observations and does not constitute reliable evidence for the prediction.

The effect was category-dependent at the descriptive level: objects showed the strongest modulation (+0.027), followed by false fonts (+0.012), while letters showed essentially no effect (-0.0004) and faces showed a reversal (-0.008). This category dependence may reflect differences in how the constraint architecture engages with pictorial versus symbolic stimuli, but the pattern requires replication in larger samples before interpretation.

Given the absence of a subject-level effect, Prediction C is not considered confirmed.

5. Discussion

5.1 Summary of findings

This study tested three pre-specified predictions derived from a constraint-architecture model of perceptual resolution (Jones, 2026) against the COGITATE open iEEG dataset. One prediction was supported, one was null, and one was inconclusive.

The supported prediction — that within-category task relevance delays perceptual resolution onset — represents a novel empirical finding. The effect was present in 73.5% of patients and statistically significant at the subject level ($p = 0.007$). It demonstrates that the observer’s task-dependent constraint state shapes the temporal dynamics of perceptual resolution in onset-responsive cortical electrodes.

5.2 What the onset latency effect reveals

The constraint-load onset delay is a specific and non-trivial finding. When a stimulus belongs to the same category as the current target but is not itself the target, the perceptual system must perform a finer discrimination — distinguishing “face that is the target” from “face that is not the target” — than when processing a stimulus from an entirely irrelevant category. On the constraint-architecture model, this finer discrimination requires resolving signal in a region of the perceptual space where constraint boundaries are closer together, and that resolution takes measurably longer.

This finding has a natural interpretation within predictive processing frameworks (Rao & Ballard, 1999; Friston, 2010; Clark, 2013): precision-weighted prediction error should be higher for stimuli that partially match the current task set, requiring more iterative processing to resolve. What the constraint-architecture model adds to this interpretation is a person-level account of *why* the precision weighting is structured as it is — namely, that the accumulated history of the observer (including the task instructions absorbed moments earlier) has configured a constraint architecture that treats near-target stimuli as requiring more careful routing than far-from-target stimuli.

Importantly, this effect operates in the *opposite* direction from a simple attentional facilitation account. Across categories, task-relevant stimuli show *faster* processing than irrelevant stimuli (consistent with attentional priming). But *within* a relevant category, the finer discrimination required for non-targets produces a *delay*. This dissociation — between-category facilitation paired with within-category constraint cost — is a signature that pure facilitation models do not naturally predict but that falls naturally out of a constraint-architecture framework.

5.3 What the null hysteresis result constrains

The failure of Prediction B is informative rather than merely negative. The constraint-architecture model predicted that accumulated constraint would persist across miniblock

transitions, producing elevated responses to newly irrelevant stimuli that decay over trials. The data show no such effect.

Three interpretations are consistent with both the null result and the broader structural framework. First, the miniblock transitions in the COGITATE paradigm are *explicit*: participants are shown new target stimuli at the start of each miniblock. This active re-instruction may clear the prior constraint state more rapidly than passive context shifts would. The model's prediction of hysteresis may hold in paradigms where constraint shifts are implicit or gradual rather than explicitly cued. Second, the constraint architecture may reset at the category level faster than five trials but persist at finer-grained levels (e.g., identity-specific priming). The present analysis, which averaged across all stimuli within a category, may have been too coarse to detect identity-level persistence. Third, the hysteresis signature may be better captured by measures other than mean high-gamma amplitude — for instance, by examining trial-by-trial variability, decoding confidence, or temporal generalization patterns.

The null result constrains the model: constraint persistence at the category level, following explicit context shifts, does not produce detectable high-gamma amplitude hysteresis at the timescale tested. This is a useful boundary condition.

5.4 Duration-tracking: an inconclusive result

Prediction C produced an inconclusive result. The overall direction was as predicted at the electrode-comparison level, but no subject-level effect was present. The absence of a subject-level effect means the comparison-level result cannot be interpreted as evidence for the prediction, given the non-independence of electrode observations within subjects.

The category-dependent pattern at the descriptive level — objects and false fonts showing stronger constraint modulation than faces and letters — may reflect genuine differences in how sustained representations are maintained across stimulus types, but this pattern requires independent replication before interpretation.

5.5 Implications for theories of consciousness

The COGITATE adversarial collaboration challenged both IIT and GNWT. The present reanalysis adds a dimension that neither theory specifically addressed: the temporal dynamics of perceptual resolution as a function of the observer's task-dependent constraint state.

Neither IIT nor GNWT specifically predicts a within-category, task-dependent onset-latency effect in sensory cortical responses of the kind tested here. IIT locates consciousness in the integrated information structure of posterior cortex. GNWT locates it in the global broadcasting dynamics of a fronto-parietal workspace. The 11 ms constraint-load effect demonstrated here operates at a level of within-category discrimination that neither theory's current formalism directly models.

This does not refute either theory. It identifies a gap: the temporal dynamics of perceptual resolution carry systematic structure related to the observer's constraint state, and any

complete theory of conscious content must account for this. The constraint-architecture model (Jones, 2026) provides one framework for doing so; other frameworks that incorporate precision-weighted prediction error or adaptive gain modulation may generate similar predictions. What the present data establish is that the phenomenon exists and is measurable.

A broader observation follows from these results. What the iEEG data record is the temporal signature of a perceptual resolution event—signal meeting constraint and producing a measurable outcome. The high-gamma onset is a third-person accessible record of that resolution. This is consistent with perceptual resolution being a structural phenomenon that writes records into neural activity. Whether the participant was conscious of the stimulus—and what that experience was like—is a separate question that the neural record does not directly answer. The present findings are consistent with the possibility that perceptual resolution and conscious experience, though intimately related, are structurally distinct: one is an event that produces measurable records, the other is the sustained state of being the system in which such events occur. This distinction, if correct, would have implications for how theories of consciousness frame their empirical targets.

5.6 Relationship to the structural account

The predictions tested here were derived from a structural account of perception as a bidirectional channel between accumulated constraint and conscious report (Jones, 2026). That account proposes that what reaches awareness is already classified, weighted, and routed by a constraint architecture operating faster than deliberation — and that this architecture imposes measurable temporal signatures on the resolution process.

The supported onset latency effect is consistent with this account: the constraint architecture shapes *when* content resolves, not only *what* resolves. The null hysteresis result constrains the model's claims about temporal persistence. The inconclusive duration-tracking result leaves open whether constraint modulation of sustained processing is a reliable phenomenon.

Crucially, these results were obtained by pre-specifying directional predictions and accepting the outcome — including one clear disconfirmation. The structural account is refined, not rescued, by these data. The model now predicts constraint-load delays at onset but does not predict category-level hysteresis following explicit context shifts. That is a more precise model than the one we started with.

5.7 Limitations

Several limitations should be noted. First, the iEEG sample consists of patients with pharmacologically resistant epilepsy. Although channels within the epileptic onset zone were excluded, the generalizability to neurotypical populations requires confirmation with non-invasive methods (MEG, EEG). Second, electrode placement was clinically determined, producing uneven coverage of cortical regions across patients. Onset-responsive electrodes were not constrained to any particular anatomical region. Third, the onset latency measure (50% of peak amplitude) is a summary statistic that may not capture the full complexity of

temporal dynamics; future analyses could examine the entire temporal profile using temporal decoding or representational similarity analysis. The direction of the Prediction A effect was stable across all four stimulus categories, and the subject-level result is robust to the choice of summary statistic (signed-rank test operates on ranks, not raw latency values), but formal sensitivity analysis across alternative onset definitions and smoothing kernels would further strengthen the finding. Fourth, the within-category comparison treats all Non-Target items as equivalent, but constraint load likely varies with visual similarity to the specific target — an identity-level analysis could reveal finer structure. Fifth, the distribution of latency differences was right-skewed at the electrode level (mean 11.3 ms, median 3.9 ms), but the subject-level distribution was more symmetric (mean +11.6 ms, median +14.1 ms, skewness 0.64), and three robust estimators (Hodges-Lehmann, trimmed mean, bootstrap CI) all converge with the parametric mean. The main finding is not driven by outlier subjects or a small number of extreme electrodes.

5.8 Future directions

The present findings motivate several extensions. The COGITATE consortium has announced upcoming releases of MEG ($n = 100$) and fMRI ($n = 118$) datasets from the same experimental paradigm. The onset latency effect identified here can be tested in MEG with whole-brain sensor coverage and a much larger healthy sample. The fMRI data can localize which cortical regions show the strongest constraint modulation, providing spatial precision complementary to the temporal precision of iEEG.

Beyond the COGITATE dataset, the constraint-architecture model generates predictions testable in other paradigms: perceptual decision-making under varying levels of within-category similarity, attentional set-shifting tasks where constraint transitions are implicit rather than explicit, and developmental or clinical populations where the accumulated constraint architecture may differ systematically from neurotypical adults.

6. Conclusion

We tested three predictions derived from a constraint-architecture model of perceptual resolution against the COGITATE open iEEG dataset. The central finding is that within-category task relevance delays the onset of perceptual resolution by approximately 11 milliseconds — a constraint-load effect that is consistent across stimulus categories, present in the majority of patients, and statistically significant at the subject level. Neither IIT nor GNWT specifically predicts a within-category, task-dependent onset-latency effect in sensory cortical responses of this kind, and the finding demonstrates that the observer's task-dependent constraint state shapes the temporal dynamics of perceptual resolution in measurable ways.

One prediction (category-level hysteresis) was clearly disconfirmed, constraining the model's claims about the timescale and conditions of constraint persistence. One prediction (constraint-modulated duration-tracking) produced inconclusive results.

These results contribute to the COGITATE consortium's call for new quantitative frameworks for testing theories of consciousness. They suggest that the temporal dynamics of perceptual resolution — not just its spatial distribution or global broadcasting properties — carry systematic information about the observer's constraint state, and that models of conscious content should account for this dimension.

References

- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.
- Cogitate Consortium, Ferrante, O., Gorska-Klimowska, U., et al. (2025). Adversarial testing of global neuronal workspace and integrated information theories of consciousness. *Nature*, 642(8066), 133–142.
- Dehaene, S., & Changeux, J.-P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron*, 70(2), 200–227.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
- Gramfort, A., Luessi, M., Larson, E., et al. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, 7, 267.
- Jones, J. C. (2026). The self the ego did not build: What decides before you decide. HoldingLight LLC. *PhilArchive*.
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87.
- Seedat, A., Lepauvre, A., Jeschke, J., et al. (2025). Open multi-center intracranial electroencephalography dataset with task probing conscious visual perception. *Scientific Data*, 12, 854.
- Tononi, G., Boly, M., Massimini, M., & Koch, C. (2016). Integrated information theory: From consciousness to its physical substrate. *Nature Reviews Neuroscience*, 17(7), 450–461.
-

Data and Code Availability

The COGITATE iEEG dataset (v1.2) is available at <https://www.arc-cogitate.com/data-release> under CC BY 4.0 (Seedat et al., 2025). Analysis code for all predictions is available at <https://github.com/Holding-Light/cogitate-constraint-reanalysis> and archived on OSF (<https://doi.org/10.17605/OSF.IO/MXYU2>).

AI Disclosure. AI tools were used to assist with manuscript preparation. The underlying theory, arguments, and interpretive claims are the author's own, and the author takes full responsibility for the content.

Citation: Jones, J. C. (2026). Constraint-Dependent Perceptual Resolution: A Pre-Specified Reanalysis of the COGITATE iEEG Dataset. HoldingLight LLC.

Manuscript version: v1.0 — 2026-03 HoldingLight LLC