International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 4, July-Aug-2026, ISSN (Online): 2395-566X

Predictive Processing and Narrative Suspense: Toward a Computational Neuroaesthetics of Plot

Yathish Kumar S
Department of Languages
Sri Sathya Sai University for Human Excellence, Kalaburagi, Karnataka, India

Abstract — This paper proposes a theoretical synthesis between the predictive processing framework in computational neuroscience — most rigorously formalized as Karl Friston's free energy principle — and the literary phenomenon of narrative suspense. We argue that suspense is not merely a felt quality of certain texts but a measurable consequence of how narrative structures modulate prediction error across a reader's generative model of unfolding events. Drawing on empirical work in neuroaesthetics, psycholinguistics, and narratology, we develop a model in which classical plot devices — foreshadowing, withheld information, dramatic irony, the ticking clock, and the twist — function as deliberate engineering of precision-weighted prediction error. We review existing fMRI and behavioral evidence on narrative transportation and the Default Mode Network, and we present, as a Stage 1 Registered Report, a fully specified 2×2 within-subjects experimental design with power analysis, a pre-registered statistical analysis plan using linear mixed-effects modeling, and directional hypotheses derived independently from competing theoretical accounts. We close by addressing the principal objection to this research program — that formalizing suspense in information-theoretic terms risks reducing literary value to a computable quantity — and argue that this concern rests on a confusion between mechanism and meaning.

Keywords— predictive processing, free energy principle, narrative suspense, neuroaesthetics, narratology, prediction error, Registered Report

I. INTRODUCTION

Suspense is among the oldest and least theorized effects in narrative art. Aristotle's Poetics gestures toward it through the concepts of peripeteia (reversal) and anagnorisis (recognition), but offers no causal account of why these structures produce the felt tension readers and audiences describe. Twentieth-century narratology — Genette, Bal, Sternberg — refined the formal vocabulary considerably, distinguishing suspense from curiosity and surprise along axes of temporal knowledge asymmetry between reader, narrator, and character. Yet narratology has remained almost entirely a descriptive science: it tells us what structures produce suspense without explaining why a structure of that kind should produce a state with measurable physiological correlates, including elevated heart rate, skin conductance change, and sustained attentional capture (Lehne & Koelsch, 2015).

The predictive processing (PP) framework offers, for the first time, a mechanistic bridge between this formal description and an explanatory account. Originating in computational models of perception (Rao & Ballard, 1999) and generalized into a unified theory of brain function by Friston (2010) under the free energy principle, PP holds that the brain is fundamentally a hierarchical inference engine: it continuously generates predictions about incoming sensory and cognitive data and updates its internal models in proportion to prediction error, weighted by the estimated precision, or reliability, of that error. This framework was developed to explain perception, motor control, and even psychiatric disorders, not literature. Its extension to narrative is therefore not a metaphorical borrowing but a substantive empirical claim: that the same computational architecture governing visual scene prediction also governs a reader's moment-to-moment expectations about plot.

This paper makes four contributions. First, it formalizes the analogy between narrative structure and hierarchical predictive inference with more precision than prior neuroaesthetic accounts typically offer (e.g., Hsu et al., 2014; Armstrong, 2020). Second, it maps specific literary devices onto specific operations within the PP architecture, distinguishing devices that manipulate the magnitude of prediction error from those that manipulate its precision-weighting.

Third, it specifies a complete experimental design, including a power analysis and a pre-registered statistical analysis plan, that would allow this framework to be tested directly against competing accounts. Fourth, by adopting the Registered Report format,[1] the paper distinguishes clearly between what is theoretically argued, what is empirically already known, and what remains to be tested — a distinction that has not, to date, been rigorously maintained in neuroaesthetic writing on narrative.

II. THE FREE ENERGY PRINCIPLE AND PREDICTIVE PROCESSING: A PRIMER

Under the free energy principle, biological systems persist by minimizing variational free energy, an upper bound on surprise, formally the negative log-evidence of sensory data given a generative model (Friston, 2010). The brain achieves this minimization through two complementary routes: updating internal predictions to better match incoming data, known as perceptual inference, or acting on the world so that incoming data come to match predictions, known as active inference. Prediction error, the mismatch between expected and actual input, is the central computational currency of this account. Critically, not all prediction error is treated equally: the brain assigns precision weights to error signals based on their estimated reliability, meaning that an unexpected event occurring in a context of high uncertainty produces a different inferential response than the identical event occurring in a context of confident expectation (Clark, 2013).

This architecture is hierarchical. Higher cortical levels generate predictions about more abstract, slower-changing features of the environment — for instance, that a text identified as a detective story likely concerns a concealed crime — while lower levels handle finer-grained, faster-changing predictions, such as the next word in a sentence. Prediction errors that cannot be resolved at a lower level propagate upward, forcing revision of higher-level models, a process directly analogous to what narratologists describe as re-reading a plot in light of new information.

The relevance to narrative becomes immediate once reading itself is recognized as a continuous act of generative modeling. A reader does not passively receive a sequence of events; she constructs, updates, and revises a model of the storyworld, including its causal structure, character motivations, and likely trajectory, at every level from syntax to plot (Zwaan, 2004; Kuijpers et al., 2014, on narrative absorption).

III. SUSPENSE AS PRECISION-WEIGHTED PREDICTION ERROR

The central theoretical claim of this paper is that narrative suspense is the subjective correlate of sustained, high-precision prediction error that the reader's generative model cannot yet resolve.

This formulation distinguishes suspense from two adjacent but distinct narrative affects. Surprise corresponds to a single, large-magnitude prediction error event accompanied by low anticipated precision beforehand, as when a reader did not expect to need a prediction at all, for example an unforeshadowed plot twist. Curiosity corresponds to prediction error in the absence of strong directional expectation, an open epistemic gap populated by multiple low-precision hypotheses, as in the question of a murderer's identity before any suspects have been introduced. Suspense, by contrast, requires the text first to elevate precision, making the reader confident that something specific and consequential is imminent, and then to delay resolution.

Sternberg's (1978) classic narratological distinction between suspense and curiosity maps with striking cleanness onto this precision-versus-magnitude distinction, but predictive processing supplies the causal mechanism that narratology itself cannot: the felt urgency of suspense is not merely an effect of knowing that something bad might happen, but a consequence of the brain's allocation of attentional and metabolic resources toward resolving a high-precision, unresolved error, a state that the free energy account predicts should be measurably costly (Friston, 2010; Pezzulo et al., 2015, on the link between precision-weighting and effortful attention).

1. A Taxonomy of Suspense-Generating Devices

Classical narrative devices can be classified according to which parameter of the predictive processing architecture they manipulate, as summarized in Table 1.

Device	PP Operation	Mechanism
Foreshadowing	Raises precision of a future prediction	Cues the reader's generative model to assign high confidence to a specific future event class
Dramatic irony	Creates precision asymmetry between reader- and character-level models	Reader's model is more precise than the character's, generating sustained error around when the character's model will update
The ticking clock	Adds a temporal decay function to unresolved precision	Converts an open prediction into one with increasing stakes per unit of narrative time

Withheld information	Suppresses resolution despite high precision	Prevents the error-minimizing update the system is primed to seek
The twist	Retroactively reveals that precision was misallocated	Forces wholesale revision of the higher-level generative model rather than a routine update

This taxonomy yields a non-trivial, testable prediction: foreshadowing and dramatic irony should be neurally dissociable from twists, because the former operate by maintaining an active, unresolved prediction, generating sustained activity in error-monitoring circuitry, while the latter operate through abrupt model revision, a signature more consistent with transient, large-amplitude prediction error responses, comparable to P300-like components observed in studies of narrative event boundaries (Bezdek & Schumacher, 2017).

IV. EXISTING EMPIRICAL EVIDENCE AND ITS LIMITS

1. Narrative Transportation and the Default Mode Network

fMRI studies consistently show that immersive narrative engagement recruits regions associated with the Default Mode Network, including medial prefrontal cortex, posterior cingulate, and temporoparietal junction, alongside theory-of-mind and episodic simulation networks (Mar, 2011; Hsu et al., 2014). This pattern is consistent with the predictive processing account of narrative comprehension as active generative simulation of a storyworld model rather than passive decoding of sequential information.

2. Suspense and Physiological Arousal

Lehne and Koelsch's (2015) review of suspense across music and film reports consistent findings of elevated skin conductance and heart rate deceleration during high-suspense passages, with arousal tracking the anticipated stakes of resolution rather than resolution itself, a pattern directly predicted by a precision-weighted account, since arousal should track the model's confidence in an impending high-stakes prediction error rather than the error itself.

3. Prediction Error in Language Comprehension

At the sentence level, N400 and P600 event-related potential components index prediction error during real-time language processing (Kuperberg & Jaeger, 2016), providing a well-established lower-level analogue for the higher-level, plot-scale predictions addressed in this paper. The open empirical question, and a clear gap in the current literature, is whether this lower-level predictive machinery shares computational resources with plot-level narrative prediction, or whether the latter constitutes a functionally distinct hierarchical level recruiting different neural populations.

4. The Critical Missing Experiment

No study to date has directly manipulated precision-weighting at the plot level, for example by experimentally varying the confidence with which a foreshadowed event is cued independent of its narrative content, while simultaneously measuring neural or physiological response. Sections 5 and 6 specify such a study in full, including statistical power and a pre-registered analysis plan.

V. METHOD

This section is written in the Registered Report format. All hypotheses, design parameters, and analysis procedures are specified prior to data collection.

1. Participants

We propose recruiting N = 64 healthy adult native speakers of the stimulus language, aged 18–45, with normal or corrected-to-normal vision, no history of neurological or psychiatric disorder, and no current use of psychoactive medication affecting autonomic function. Sample size is determined by the power analysis reported in Section 5.5. Participants will be screened for genre familiarity with detective and thriller fiction using a short forced-choice familiarity inventory, and genre familiarity will be entered as a covariate (see Section 5.6) given its predicted moderating role (Section 6, Hypothesis 4).

2. Design

The study employs a 2 (Precision: high vs. low) × 2 (Resolution Delay: short vs. long) within-subjects factorial design, fully crossed, with stimulus item as an additional random factor. Each participant reads 80 short narrative vignettes (20 per cell), each 150–220 words, constructed so that a target outcome event is cued and then resolved.

- High precision: the vignette contains an explicit, unambiguous cue that a specific, named outcome is forthcoming (e.g., a character explicitly states what will happen).
- Low precision: the vignette contains a comparable cue, matched for valence and stakes, but ambiguous as to the specific outcome (multiple plausible outcomes remain live).

- Short delay: 1–2 sentences (approx. 15–25 words) separate cue from resolution.
- Long delay: 5–7 sentences (approx. 80–110 words) separate cue from resolution, with narrative content held neutral with respect to the target prediction (filler description, not new information).

Item sets will be normed in a separate pilot sample (n = 30, not included in the main analysis) for perceived stakes, valence, and comprehensibility, and matched across cells on these dimensions using one-way ANOVA with a pre-specified equivalence threshold of $|d| < 0.2$.

3. Stimuli Construction and Norming

All vignettes will be authored by the research team and reviewed by two independent raters blind to condition for (a) the intended precision manipulation, verified via a forced-choice prediction task in the norming sample, and (b) absence of confounding lexical cues (sentence length, word frequency, syntactic complexity) using corpus-based readability metrics (Flesch-Kincaid, word frequency via SUBTLEX norms). Any item pair differing by more than one grade level in readability between conditions will be revised or excluded.

4. Procedure and Measures

Participants will read vignettes self-paced, word-by-word or phrase-by-phrase (RSVP), while continuous skin conductance and pupillometry are recorded; a subset of participants (n = 32, power permitting per Section 5.5) will additionally undergo fMRI using an event-related design with jittered stimulus onset asynchronies (4–8 s) to permit deconvolution of cue and resolution-locked hemodynamic responses. Primary dependent measures are:

- Skin conductance response (SCR) amplitude in the cue-to-resolution window, baseline-corrected.
- Pupil dilation, peak amplitude and time-to-peak, in the same window.
- Self-reported suspense rating (0–100 visual analogue scale) immediately following each vignette's resolution.
- (fMRI subsample only) BOLD signal in anterior insula and dorsal anterior cingulate cortex, regions selected a priori for their established role in interoceptive prediction error (Seth & Friston, 2016).
- Reading time at the resolution sentence, as an index of processing cost at the point of prediction confirmation or revision.

5. Power Analysis

Sample size was determined via Monte Carlo simulation for a linear mixed-effects model with crossed random effects for participants and items (Westfall, Kenny, & Judd, 2014), the appropriate approach for designs with repeated stimulus sampling. Based on the smallest comparable effect size reported in prior suspense psychophysiology research (Lehne & Koelsch, 2015, Cohen's $d \approx 0.45$ for high vs. low tension conditions on skin conductance), and assuming a conservative reduction to $d = 0.35$ for the more tightly matched precision manipulation proposed here, simulations across 1,000 iterations indicate that N = 64 participants with 20 items per cell provides 90% power to detect the predicted Precision × Delay interaction ($\alpha = .05$, two-tailed) on the primary SCR measure. The fMRI subsample (n = 32) is powered at 80% to detect a medium region-of-interest effect ($d = 0.5$) based on prior interoceptive prediction error studies, consistent with recommended minimum samples for region-of-interest fMRI designs (Poldrack et al., 2017). Full simulation code and assumptions will be deposited on the Open Science Framework prior to data collection.

6. Statistical Analysis Plan

All primary analyses will use linear mixed-effects models (lme4 in R), with Precision (high/low), Delay (short/long), and their interaction as fixed effects, and random intercepts and slopes for participants and items (maximal model justified by design, Barr et al., 2013). Genre familiarity (centered, continuous) will be entered as a fixed covariate, with its interaction with Precision retained only if it improves model fit by likelihood ratio test ($\Delta AIC > 2$).

The pre-registered model specification, in R-style notation, is:
SCR_amplitude ~ Precision * Delay + GenreFamiliarity + (1 + Precision * Delay | Participant) + (1 | Item)

Identical model structure will be applied to pupil dilation, suspense ratings (ordinal mixed model, cumulative link), reading time at resolution (log-transformed), and, for the fMRI subsample, region-of-interest beta estimates extracted from anterior insula and dACC. Whole-brain analyses will additionally be conducted at a cluster-forming threshold of $p < .001$ with cluster-level family-wise error correction at $p < .05$ (Eklund, Nichols, & Knutsson, 2016, informing the choice of a conservative threshold given known inflation risks in fMRI cluster inference).

Significance for all primary confirmatory tests is set at $\alpha = .05$, two-tailed, with Holm-Bonferroni correction applied across the five primary dependent measures (corrected family-wise $\alpha = .05$). Exploratory analyses, including individual-difference moderation by trait absorption (using the Tellegen Absorption Scale) and by self-reported reading frequency, will be clearly labeled as exploratory and will not be used to qualify confirmatory conclusions.

VI. HYPOTHESES AND PREDICTED OUTCOMES

The following hypotheses are derived independently from the predictive processing account and are stated so as to be falsifiable against at least one plausible rival account.

- H1 (Interaction hypothesis). SCR amplitude and pupil dilation in the cue-to-resolution window will show a significant Precision × Delay interaction, such that the high-precision, long-delay cell produces disproportionately greater arousal than predicted by the additive combination of either factor alone. A purely additive pattern (no interaction) would favor a generic "tension accumulation" account over the precision-weighted account proposed here.
- H2 (Neural dissociation hypothesis). In the fMRI subsample, anterior insula and dACC activation will scale with the high-precision/long-delay condition more than an additive model predicts, while activation at the moment of an unforeshadowed twist-like resolution (a planned exploratory item subset) will instead show a transient, large-amplitude response without preceding sustained elevation, consistent with the mechanistic dissociation proposed in Section 3.1.
- H3 (Resolution dynamics hypothesis). At resolution, the high-precision condition will show a sharper, more transient psychophysiological response than the low-precision condition, which is predicted instead to show a broader, more diffuse response consistent with parallel hypothesis elimination under a curiosity-based account (Sternberg, 1978).
- H4 (Moderation hypothesis). Genre familiarity will moderate the Precision effect, such that genre-familiar readers show steeper precision-driven arousal gradients, consistent with the view that generative models are calibrated by prior exposure to genre convention.

Disconfirmation conditions are specified as follows: if no Precision × Delay interaction emerges (H1 null) and effects are purely additive or driven by Delay alone, this would constitute evidence against the precision-weighted account in favor of a simpler arousal-accumulation model, and would be reported as such regardless of outcome.

Objections and Limitations

The Reductionism Objection

A natural and serious objection holds that modeling suspense as precision-weighted prediction error reduces a richly meaningful aesthetic experience to a computational quantity, draining literary works of the interpretive depth that is their actual scholarly interest. This paper takes the objection seriously but distinguishes two claims that are often conflated: a mechanistic claim about how suspense is computationally instantiated, and an evaluative claim about what suspense means or is worth in a given text. The predictive processing account is offered only as the former. It explains why a structurally similar device, such as a ticking clock, produces a suspense-like state across radically different texts, without claiming to explain, or replace, what any particular deployment of that device means within, for instance, a postcolonial detective novel as opposed to a Cold War thriller. Mechanism and meaning are not competitors; cognitive science has long maintained comparable distinctions in vision science, where explaining edge detection mechanistically does not exhaust the meaning of a painting.

The Individual Differences Problem

Generative models are individually calibrated through prior experience with genre conventions, such that a reader unfamiliar with detective fiction will hold different baseline precision estimates than a genre-literate reader encountering the same foreshadowing cue. This predicts measurable individual variation in suspense response correlated with genre familiarity, itself addressed as Hypothesis 4, but it complicates any simple stimulus-response account and must be controlled for carefully in both stimulus norming and analysis.

The Ecological Validity Problem

Laboratory narrative stimuli, necessarily short and decontextualized, may not recruit the same hierarchical depth of prediction as a four-hundred-page novel read over weeks, in which higher-level predictions about character, theme, and authorial intention operate on a timescale current neuroimaging methods cannot easily capture. This is a genuine limitation of the proposed paradigm, openly acknowledged here, and an open methodological challenge for the field as a whole rather than a flaw specific to this design.

VII. CONCLUSION

The free energy principle was not designed with literature in mind, yet it offers literary studies something narratology alone has never had: a causal, falsifiable, cross-domain account of why certain structures reliably produce the experience readers call suspense. This paper has argued that classical narrative devices map onto distinct, identifiable operations within a hierarchical predictive architecture, that existing, if indirect, neuroscientific evidence is broadly consistent with this account, and it has specified a complete experimental and statistical protocol, in Registered Report format, capable of testing this account directly against rival theories before a single participant is run. Whether or not the resulting data vindicate the strong form of this hypothesis, the act of specifying it in advance, with power, design, and analysis fixed, is itself the methodological contribution: it moves the cognitive study of narrative form from suggestive analogy toward a discipline answerable to its own predictions.

REFERENCES

1. Armstrong, P. B. (2020). Stories and the brain: The neuroscience of narrative. Johns Hopkins University Press.
2. Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. Journal of Memory and Language, 68(3), 255–278.
3. Bezdek, M. A., & Schumacher, E. H. (2017). Cinematic shifts in attention: Event-related potential evidence of narrative engagement. Journal of Cognitive Neuroscience.
4. Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behavioral and Brain Sciences, 36(3), 181–204.
5. Eklund, A., Nichols, T. E., & Knutsson, H. (2016). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. Proceedings of the National Academy of Sciences, 113(28), 7900–7905.
6. Friston, K. (2010). The free-energy principle: A unified brain theory? Nature Reviews Neuroscience, 11(2), 127–138.
7. Hsu, C.-T., Conrad, M., & Jacobs, A. M. (2014). Fiction feelings in Harry Potter: Haemodynamic response in the mid-cingulate cortex correlates with immersive reading experience. NeuroReport.
8. Kuijpers, M. M., Hakemulder, F., Tan, E. S., & Doicaru, M. M. (2014). Exploring absorbing reading experiences. Scientific Study of Literature.
9. Kuperberg, G. R., & Jaeger, T. F. (2016). What do we mean by prediction in language comprehension? Language, Cognition and Neuroscience.
10. Lehne, M., & Koelsch, S. (2015). Toward a general psychological model of tension and suspense. Frontiers in Psychology, 6, 79.
11. Mar, R. A. (2011). The neural bases of social cognition and story comprehension. Annual Review of Psychology, 62, 103–134.
12. Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active inference, homeostatic regulation and adaptive behavioural control. Progress in Neurobiology.
13. Poldrack, R. A., Baker, C. I., Durnez, J., Gorgolewski, K. J., Matthews, P. M., Munafò, M. R., ... & Yarkoni, T. (2017). Scanning the horizon: Towards transparent and reproducible neuroimaging research. Nature Reviews Neuroscience, 18(2), 115–126.
14. 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.
15. Seth, A. K., & Friston, K. J. (2016). Active interoceptive inference and the emotional brain. Philosophical Transactions of the Royal Society B.
16. Sternberg, M. (1978). Expositional modes and temporal ordering in fiction. Johns Hopkins University Press.
17. Westfall, J., Kenny, D. A., & Judd, C. M. (2014). Statistical power and optimal design in experiments in which samples of participants respond to samples of stimuli. Journal of Experimental Psychology: General, 143(5), 2020–2045.
18. Zwaan, R. A. (2004). The immersed experiencer: Toward an embodied theory of language comprehension. Psychology of Learning and Motivation, 44, 35–62.