ChatGPT

Active Listening Beside the Record: Priors, Authorship, and the Limits of Inference

Abstract

This paper asks whether Karl Friston and colleagues' model of "active listening" can illuminate the reading of heterogeneous documentary records without turning a computational theory of speech perception into an unwarranted theory of archival truth. Friston et al. model speech recognition as inference over competing segmentations of a continuous acoustic stream: a simulated listener covertly places boundaries, evaluates alternative lexical hypotheses, and favors interpretations that maximize model evidence under prior beliefs. The paper is explicitly computational: it demonstrates face validity through simulations and proposes links between simulated belief updating and electrophysiological phenomena such as mismatch negativity, P300, and N400; it does **not** report an experiment establishing that archival readers, historians, courts, or scientific reviewers literally perform the same calculation. filecite turn3file7 [1]

Read beside the uploaded **RITR_READ_BESIDE_v1.0** and **READER R2: the COMINI and autism series**, however, active listening supplies a productive analogy with a sharp limit. RITR requires a reader to preserve speakers, documents, absences, transformations, and unresolved questions rather than collapsing them into a single explanatory voice. filecite turn2file2 R2 performs a related provenance operation: it repeatedly distinguishes what a file says from what the reader may conclude about it, beginning with the declaration that its work is read-only, has "Authority delta = 0," and treats document claims as `SOURCE_SAYS`, not findings. filecite turn3file0

The strongest discovery developed here is that **documentary reading has a segmentation problem structurally analogous to active listening, but provenance supplies constraints that semantic model evidence cannot replace**. A coherent interpretation can achieve excellent "fit" while still merging distinct speakers, mistaking an AI-generated report for an independently reviewed paper, treating a citation as authorship, or converting a speculative causal chain into an empirical result. In archival work, therefore, surprise should sometimes be conserved rather than minimized. A contradiction between metadata, byline, self-description, and later attribution is not necessarily noise to be resolved; it may be the most information-rich feature in the record.

Active listening as model, simulation, and analogy

Friston et al. begin from a genuine difficulty in speech perception: continuous acoustic input does not arrive already divided into words, and context influences lexical interpretation. Their formal solution couples recognition with segmentation. The simulated listener entertains candidate boundaries and lexical hypotheses and selects segmentations that minimize variational free energy—or, equivalently in their formulation, maximize evidence for its generative model. filecite turn2file0 [2]

Four components matter for the present inquiry.

Covert segmentation is the selection of boundaries inside already sampled sensory data. Friston et al. call the boundary placement an internal or covert action, distinguishing it from an overt action such as moving one's head or asking a question. filecite turn1file8 In document reading, an analogous operation occurs whenever a reader decides that one paragraph belongs to an author's voice, another is quoted material, another is machine-generated synthesis, and another is metadata. The analogy is useful because these divisions are not always pre-given. Yet nothing in Friston et al. demonstrates a neurological mechanism for distinguishing PDF authorship, chain of custody, quotation ownership, or AI provenance. Calling archival parsing "covert segmentation" is therefore an interpretive transfer, not a neuroscientific finding.

Context priors constrain which interpretations become probable. In Friston et al.'s simulations, changing priors or background noise changes inferred lexical content; the "Grade A"/"grey day" type of ambiguity is resolved by contextual plausibility rather than acoustic form alone. filecite turn2file0 [2] This gives the documentary analogy real force: if a reader begins with the prior "COMINI IV is Bruce's paper," then references to Bruce, Wright, COMINI, and user-supplied material can all become evidence *for* that prior unless record types are separated first. If the reader begins instead with "metadata is always dispositive," a printed human byline may be improperly discounted. Neither prior is sufficient.

Model evidence in active listening means evidence for a generative explanation of sensory data. It is not a synonym for juridical evidence, historical proof, or scientific confirmation. Friston et al. explicitly use a simplified generative model and simulations rather than a state-of-the-art speech-recognition system or a direct human-neurophysiology validation study. [2] The article describes its neuronal associations as predictions or predictive-validity considerations; the numerical relationships depicted in its Bayesian-surprise analyses are simulated. filecite turn1file7 filecite turn1file9 Treating "maximal model evidence" as though it licensed a historical authorship claim would therefore commit a category error.

Bayesian surprise is the amount of belief updating induced by new information. Friston et al.'s simulations predict larger update-related responses where accurate priors are absent; Figure 10 reports a monotonic relation between simulated Bayesian surprise and simulated evoked activity. filecite turn1file7 The important word is **simulated**. The article connects the model prospectively to MMN, P300, and N400 phenomena, but it does not establish that a document reader's felt surprise, doubt, or revision is an EEG-indexed measure of archival truth. [3]

This distinction provides the first methodological boundary of the present paper:

> Active listening can illuminate how interpretations become selected; it cannot, by itself, determine which documentary attribution is true.

That limit is not a weakness in Friston et al. It is what allows the model to remain scientifically specific.

Reading beside rather than reading through

RITR's central discipline is almost the inverse of an over-aggressive predictive reader. It instructs the writer to "read actual passages," "preserve authorship," distinguish removal from absence, and leave unresolved relations unresolved rather than manufacturing a cleaner narrative. filecite turn2file2 It further warns that an excluded phrase, a noticed instance, and an observed listener effect are "three records"; that quoted

first-person statements remain with whoever actually wrote them; and that a combination merging speakers, packets, record types, or temporal relations should be named as a declined merge rather than smoothed into prose. filecite turn2file2

This creates an important tension with the intuitive lesson one might take from predictive processing. If a generative system prefers interpretations that make the incoming stream unsurprising, then a careless archival reader may behave as an overconfident predictor: she makes the record easier to explain by moving boundaries until everything belongs to one story.

RITR instead demands **boundary resistance**.

That resistance is visible in its treatment of Bruce's sentence, "every thing you say is a problem is actually a feature in this particular scientific work." RITR does not convert that sentence into proof that every apparent conflict is resolved. It says a feature is an occurrence that may be placed beside another occurrence; the juxtaposition does not establish that both statements are true or that no conflict exists. filecite turn2file2 This is epistemically significant. A generative model tries to explain observations. A documentary method sometimes has to preserve the failure of observations to collapse into one explanation.

The archival analogue of covert segmentation is therefore not merely "choose the best boundary." It is:

choose only those boundaries the record warrants, and preserve boundary uncertainty when the record does not settle it.

That is where "reading beside" differs from "reading through." Reading through makes one text explanatory of another. Reading beside allows two records to remain mutually informative without making either the master key.

R2 as a provenance stress test

R2 is unusually useful because its strongest claims are not substantive claims about longevity or autism. Its strongest claims are **claims about the status of claims**.

Its opening rule is explicit: read-only, nothing signed or sent, zero authority delta, and document descriptions treated as `SOURCE_SAYS`, not findings. filecite turn3file0 This distinction should govern every use of R2 in the present paper. I have inspected the uploaded R2 report; I have **not** inspected the unuploaded local files whose filesystem paths R2 names. Consequently, statements below about those files are statements about what R2 reports having observed, not independent confirmations by me.

R2 reports a particularly revealing authorship conflict around **COMINI IV: A Composite Framework for Longevity Prediction**. According to R2, the PDF metadata identifies "ChatGPT Deep Research" as Author, "ChatGPT" as Creator, and WeasyPrint as Producer; the body describes itself as a "deep-research report," refers to "the COMINI paper's author" and "the user" in the third person, and cites several Wright-authored source items. R2 therefore refuses to infer that Bruce authored COMINI IV merely because Bruce supplied ideas or appears in its citation chain. filecite turn3file2

This is exactly where a loose active-listening analogy would mislead. A high-level narrative—

Bruce developed COMINI → COMINI IV discusses COMINI → therefore COMINI IV is Bruce's paper

—has excellent semantic coherence. It may even minimize interpretive surprise. Yet R2 shows why coherence is insufficient. The record contains different evidentiary channels: metadata, printed byline, first-person usage, self-description, citation attribution, upload references, and an external desk's description. Those channels do not all answer the same question. filecite turn3file3

R2 finds an even more instructive second-order problem. One COMINI source reportedly prints Bruce Norman Wright as its byline while its PDF metadata identifies ChatGPT Deep Research as author; COMINI IV then reportedly cites that source as Wright's work. R2 does not resolve whether Bruce wrote, edited, commissioned, or supplied the knowledge base for the text. It records the attribution chain as conflicted. filecite turn2file1

This distinction parallels a mature understanding of digital provenance. NIST defines provenance in terms of origin, development, ownership, location, modification, personnel, and processes—not merely semantic resemblance. [4] NIST guidance for digital evidence likewise stresses documentation of original source, acquisition history, hashes, and custody because digital files are easy to alter. [5] R2's hashes and copy counts, if accepted as R2's reported observations, belong to this provenance layer; they establish or test file identity, not intellectual authorship by themselves. R2 reports, for example, multiple byte-identical copies and zero same-basename/different-byte variants for the specified set. filecite turn3file0 That is a file-integrity observation. It does not answer who originated the propositions inside the files.

The same separation is crucial for the autism materials. R2 describes one artifact as a brudo-generated machine report with "anonymous" PDF authorship metadata and says the file itself does not record whether its content came from a model call or Bruce's own input. filecite turn1file1 Another report is summarized as an AI-generated claim map about an autism hypothesis; another AI-generated review is explicitly described by its own text as leaving the proposed breast-pump/autism linkage theoretical. filecite turn1file4 Those are provenance and epistemic-status observations. They do **not** establish the proposed autism causation, and this paper does not adopt it.

Likewise, R2 identifies breathing-method language and quarantined numerical claims in parts of the corpus but deliberately locates rather than republishes them. filecite turn3file6 I preserve that decision here: no regimen, target, dosage, progression, killed figure, or quarantined emergency metric is reproduced or endorsed.

A comparable problem occurs with a file styled as a "Peer Review." R2 reports that its reviewer personas are simulated, its PDF metadata identifies ChatGPT Deep Research, and it is not a review by a journal or named independent reviewer. filecite turn3file6 Calling that artifact "peer reviewed" would be an inferential boundary error: the linguistic form of peer review would have been mistaken for the institutional event of peer review.

Crossref's bibliographic guidance reinforces the general point that contributor identities, titles, dates, identifiers, and roles are distinct metadata that should be accurate rather than forced into ill-fitting fields. [6] That principle does not resolve R2's authorship dispute, but it explains why "author," "generator," "commissioner," "source author," and "subject of discussion" should not be collapsed into one role.

The strongest discovery: epistemic surprise should sometimes be preserved

The most productive result of placing Friston beside RITR and R2 is not that archival research "works like the brain." That claim would be too strong. The stronger result is methodological:

An archive becomes unreliable when the reader treats explanatory fluency as provenance.

Friston et al. show why a generative system needs priors. Without priors, ambiguous sensory data are difficult to segment; with useful priors, recognition becomes tractable. [2] The documentary reader likewise cannot work without priors. Even recognizing something as "metadata," "citation," "quotation," "AI report," or "authored manuscript" presupposes categories.

The danger begins when a prior ceases to be revisable.

Suppose the prior is H_B: "Bruce authored this document." A reader then observes Bruce's name in references, concepts associated with Bruce, first-person phrasing elsewhere in the corpus, and an earlier desk calling the file "your paper." Each may increase subjective confidence in H_B. But R2 introduces observations from another evidentiary layer: ChatGPT authorship metadata, self-description as a deep-research report, third-person reference to "the user," and absence of a Bruce byline in that document. filecite turn3file2 A well-calibrated reader should update.

Yet the update cannot be completed by choosing whichever bundle "fits" best. The variables are underdetermined. A human can author ideas, prompts, passages, revisions, or an underlying knowledge base without being the textual generator of a particular PDF. A machine-generated report can cite a human as an author of antecedent material. A byline can disagree with file metadata. Commissioning is not authorship; metadata authorship is not necessarily legal or scholarly authorship; citation is not provenance; and textual voice is not a cryptographic signature.

The appropriate posterior is therefore not a crisp authorship assignment but a **structured unresolved state**.

That is the paper's central contribution: RITR's refusal to merge record types can be interpreted as an epistemic safeguard against premature posterior collapse.

This also changes how "Bayesian surprise" should be used metaphorically. In Friston et al., surprise drives belief updating. filecite turn1file9 In archival research, a contradiction can do more: it can signal that the hypothesis space is wrongly specified. If "Bruce-authored" and "AI-authored" are treated as exhaustive mutually exclusive categories, the reader may miss intermediate production relations such as commissioned generation, edited generation, source authorship, compilation, or mixed human-machine production.

Thus some archival surprises should not immediately be explained away. They should be **preserved as model criticism**.

The no-break-space observation in R2 is a small but instructive case. R2 reports that the COMINI IV filename contains U+00A0 where another source renders an ordinary space, enough to make a literal script fail. filecite turn3file0 Semantically, the two names “look the same.” Computationally, they are not the same string. The episode condenses the paper’s thesis: what appears equivalent under one model can be distinct under another, and the distinction can matter operationally.

The same logic applies to scientific hypotheses. A report can present a causal mechanism coherently, cite adjacent physiological literatures, and still lack direct evidence for the proposed causal link. R2’s own language repeatedly guards this boundary. One autism-related review is described as stating that the linkage “remains theoretical”; another is marked as a hypothesis whose contentious claims must not be adopted by the cover sheet. filecite turn1file4 Scientific evaluation should therefore preserve at least four levels:

1. **Source statement:** what a document asserts.
2. **Simulation/model result:** what follows inside a stipulated computational model.
3. **Inference:** what a reader reasonably derives from multiple records.
4. **Established empirical finding:** what has been directly observed under an empirical research design.

Friston et al.’s numerical speech-recognition and neural-response demonstrations belong primarily to level 2, while their discussion connects them to established experimental literatures at level 4 and proposes predictive links requiring further testing. [2] R2’s descriptions of COMINI and autism reports belong primarily to level 1, with its metadata, hash, and corpus observations presented as its own documented checks. The present paper’s correspondence between active listening and archival reading belongs to level 3. Keeping those levels distinct is the core protection against converting analogy into mechanism or hypothesis into fact.

The unresolved question is therefore more interesting than “Who really authored COMINI IV?” It is:

What documentary evidence would justify moving from a provenance graph containing generator, commissioner, source author, editor, and byline to one scholarly authorship statement?

The supplied record does not settle that question. R2 explicitly leaves it for Bruce to decide whether he asserts authorship of particular texts and on what basis. filecite turn3file5 A future evidentiary answer could include version history, source drafts, editing records, contemporaneous prompts, signed attestations, or production logs. None of those unuploaded materials is claimed here to have been inspected.

Editorial accounting under RITR Read Beside

Reading extent. I read the complete supplied RITR text; the uploaded R2 report’s summary, scope statement, COMINI and autism file descriptions, authorship-conflict section, default-status section, and SCOUT corrections retrieved from the document; and the Friston article’s abstract, introduction, formal segmentation discussion, Bayesian-surprise discussion, Figure 10 explanation, and discussion, with the peer-reviewed article cross-checked against its PubMed/PMC record. filecite turn2file2 filecite turn3file0 filecite turn3file2 filecite turn3file8 filecite turn2file0 [1] I did

not inspect the local `/Users/bnw/...` documents named inside R2 because they were not among the three uploaded sources.

Content grade, before: A−. This is my grade of the source constellation as material for the inquiry, not a claim that any source is deficient. R2 supplies unusually strong provenance distinctions—especially its opening rule that document claims are `SOURCE_SAYS`, "not a finding." filecite turn3file0 Friston supplies a rigorous computational account of segmentation and prior-dependent inference. filecite turn3file7 What is absent before the present paper is an explicit theory of why the two belong beside one another without claiming neurological identity.

Content grade, after: A. The paper adds that missing relation: explanatory fit and provenance are different evidentiary currencies. The central addition is the proposition that RITR's declined merges function as protection against **premature posterior collapse** in documentary interpretation. That concept is introduced here as an analogy/inference, not attributed to Friston, RITR, or R2.

Signal grade, before: B+. Again this is my editorial judgment, not a correction imposed on the source. RITR has very high procedural signal but deliberately preserves dense distinctions; R2 contains substantial operational detail, hashes, paths, and file inventories whose purpose is documentary rather than argumentative. Its clearest signal is the short sentence "Authority delta = 0." filecite turn3file0

Signal grade, after: A. The present paper concentrates the inquiry around one governing distinction: **a model can explain a record without proving its provenance**. That sentence provides a reusable test for authorship, peer-review status, scientific causation, and simulation claims while leaving the underlying unresolved records intact.

Clearest addition. The strongest new contribution is the distinction between *surprise to be minimized* and *surprise to be conserved*. Friston's simulated listener uses surprise reduction to recognize speech; the archival reader sometimes preserves contradiction because contradiction can reveal an underspecified hypothesis space. This is an original inference from placing the works beside one another, not a claim found in any source.

Clearest removal / declined merge. I declined the merge "Bruce → COMINI ideas → COMINI IV → Bruce-authored COMINI IV." R2 itself reports incompatible provenance signals and leaves authorship unresolved. filecite turn3file3 I also declined the merge "AI-generated Peer Review → independent peer review," because R2 reports simulated reviewer personas rather than an independent journal or human review event. filecite turn3file6 No autism-causation hypothesis, quarantined metric, or breathing regimen was imported into the paper as a finding.

What I remain wondering. The hardest unresolved issue is not whether priors affect interpretation—they plainly do in Friston's model—but how a rigorous documentary workflow should assign evidentiary weight when semantic authorship, textual generation, metadata authorship, commissioning, and intellectual contribution diverge. R2 identifies the divergence but appropriately does not supply a universal rule for resolving it. filecite turn3file5

EARNINGS. The named EARNINGS steps are supplied by RITR itself; I did not call or invent a separate EARNINGS tool. filecite turn2file2 **Readback:** Friston models prior-dependent covert segmentation;

RITR requires preservation of speakers and record types; R2 preserves disputed authorship rather than resolving it. **Learned:** coherent interpretation and provenance verification solve different problems. **Whydoit:** without that distinction, the reader can mistake narrative fit for authorship, simulated peer review for institutional peer review, or mechanistic plausibility for observed causation. **Giveback:** the paper returns a reusable distinction—*semantic model evidence is not provenance evidence*. **Spend:** that earning is spent in the paper's strongest methodological move: preserve certain contradictions as evidence that the current hypothesis space is inadequate rather than forcing them into the interpretation with the highest narrative coherence.

References

Friston, K. J., Sajid, N., Quiroga-Martinez, D. R., Parr, T., Price, C. J., & Holmes, E. "Active listening." *Hearing Research* 399 (2021): 107998. DOI and bibliographic record verified through PubMed/PMC. [1] Uploaded source: filecite turn2file0

RITR_READ_BESIDE_v1.0. Dated October 26, 2026; identifies Grok as author of the merge and preserves Bruce's quoted sentence without adoption into Canon. filecite turn2file2

READER R2: the COMINI and autism series. Volume 3 desk / brudo Claude desk report, dated September 26, 2026. Its claims about unuploaded local files are treated here as R2's reported observations rather than independently verified facts. filecite turn3file0

National Institute of Standards and Technology. "Provenance," CSRC Glossary; and *Digital Evidence Preservation: Considerations for Evidence Handlers*, NIST IR 8387, September 2022. [7]

Crossref. "Bibliographic metadata," contributor and metadata best-practice guidance. [6]

---

[1] Active listening
https://pubmed.ncbi.nlm.nih.gov/32732017/?utm_source=chatgpt.com

[2] Active listening
https://pmc.ncbi.nlm.nih.gov/articles/PMC7812378/?utm_source=chatgpt.com

[3] Active listening
https://www.sciencedirect.com/science/article/pii/S0378595519303491?utm_source=chatgpt.com

[4] [7] provenance - Glossary | CSRC
https://csrc.nist.gov/glossary/term/provenance?utm_source=chatgpt.com

[5] Digital Evidence Preservation: Considerations for Evidence Handlers
https://nvlpubs.nist.gov/nistpubs/ir/2022/NIST.IR.8387.pdf

[6] Bibliographic metadata - Crossref
https://www.crossref.org/documentation/principles-practices/best-practices/bibliographic/?utm_source=chatgpt.com