Quantifying information gain in narrative stimuli using open-source large language models

Goeschel, A.[1,*], Palsule, R.A.[2,*], Yoo, J.[3], Chen, J.[4], Bornstein, A.M.[2,+], Khoudary, A.[2,+]

[1] Department of Computer Science, University of California, Irvine; [2] Department of Cognitive Sciences, University of California, Irvine; [3] Princeton Neuroscience Institute, Princeton University; [4] Department of Psychological & Brain Sciences, Johns Hopkins University

[*+] These authors contributed equally.

Background

- Causal structure shapes memory encoding and retrieval (Lee & Chen, 2022; Antony et al., 2024)
- Narrative stimuli (movies, books, etc.) are powerful tools for probing effects of causal structure on cognition and behavior (Chen & Bornstein, 2024)
- Causal inference plays a key role in credit assignment to causally relevant events in long-term episodic memory (DuBrow et al., 2017; Yoo, 2025)
- Chen & Bornstein (2024) hypothesized that humans encode narratives into schema to facilitate efficient credit assignment in complex, high-dimensional environments
- We leverage recent advances in LLMs to develop a metric quantifying contextual information gain of individual narrative events: *Sequentiality*
- We evaluate what this reformulated Sequentiality captures, hypothesizing it deviates from its schema-based origins toward one more sensitive to the structure of presented information

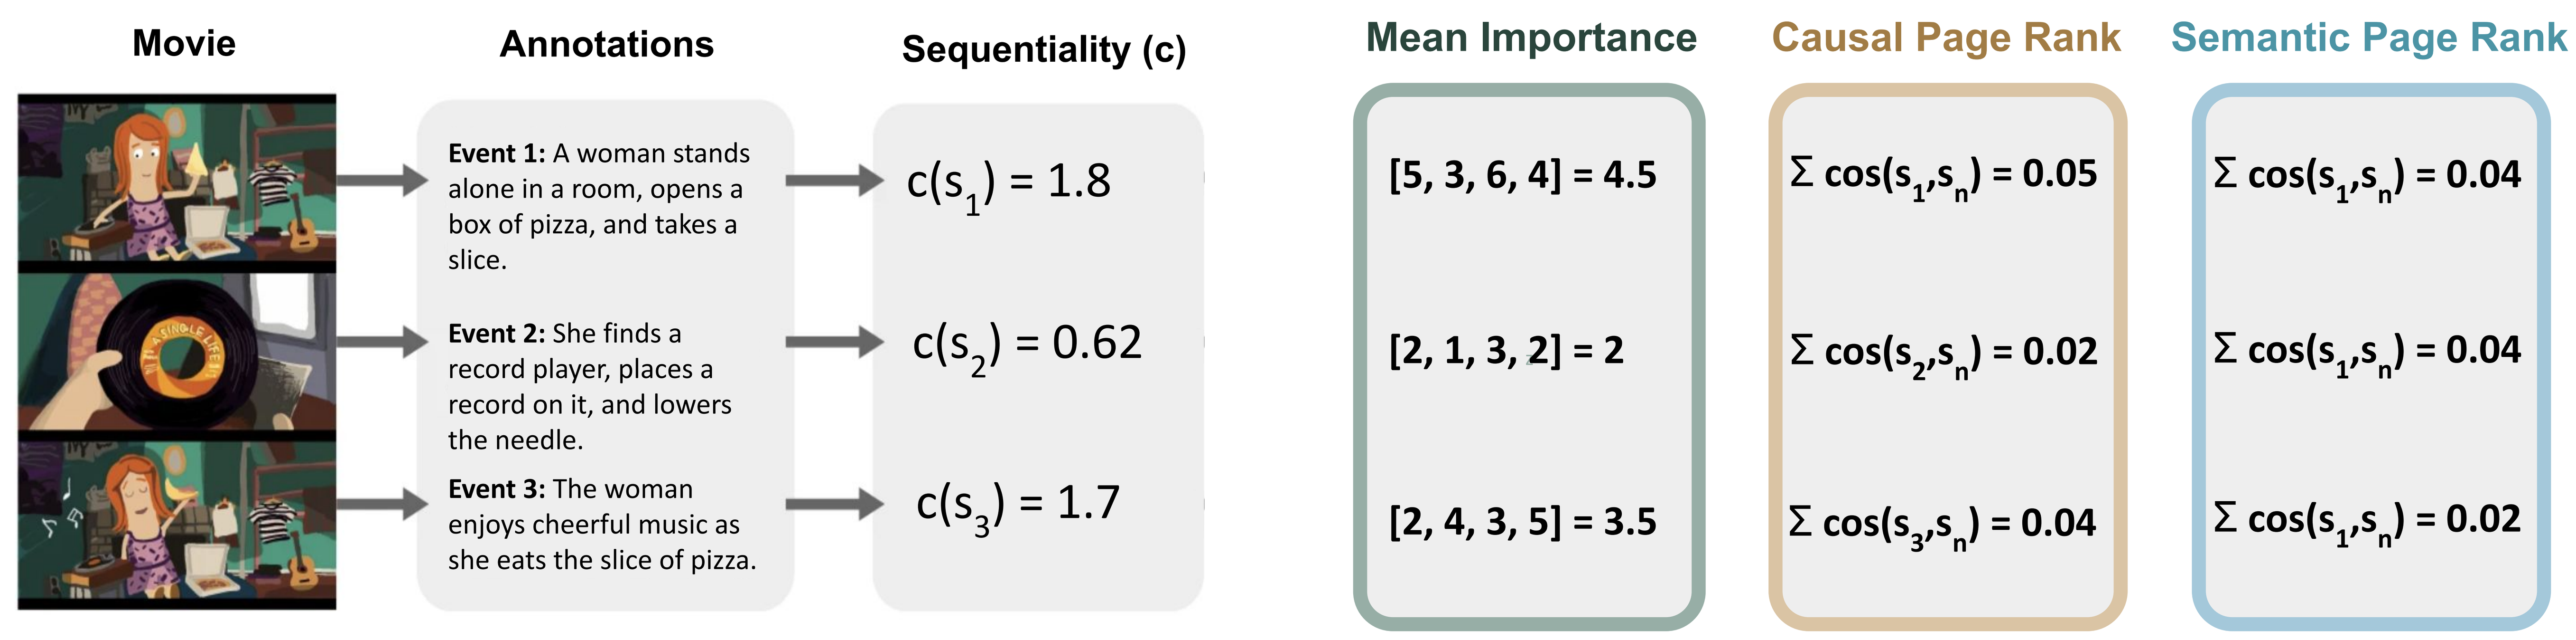
Methods: Using the FilmFest Dataset (Lee & Chen, 2022)

FilmFest Stimuli

- 10 short films, viewed in succession during fMRI scanning (n=21)
- Human-segmented event annotations (Lee & Chen, 2022) used to calculate Sequentiality values (n=216 annotations)

Event-Level Measures

- Importance: independent human ratings (n=4)
- Causal centrality: page rank over cause-effect ratings (n=12; Lee & Chen, 2022)
- Semantic similarity: page rank over USE embeddings (Lee & Chen, 2022)

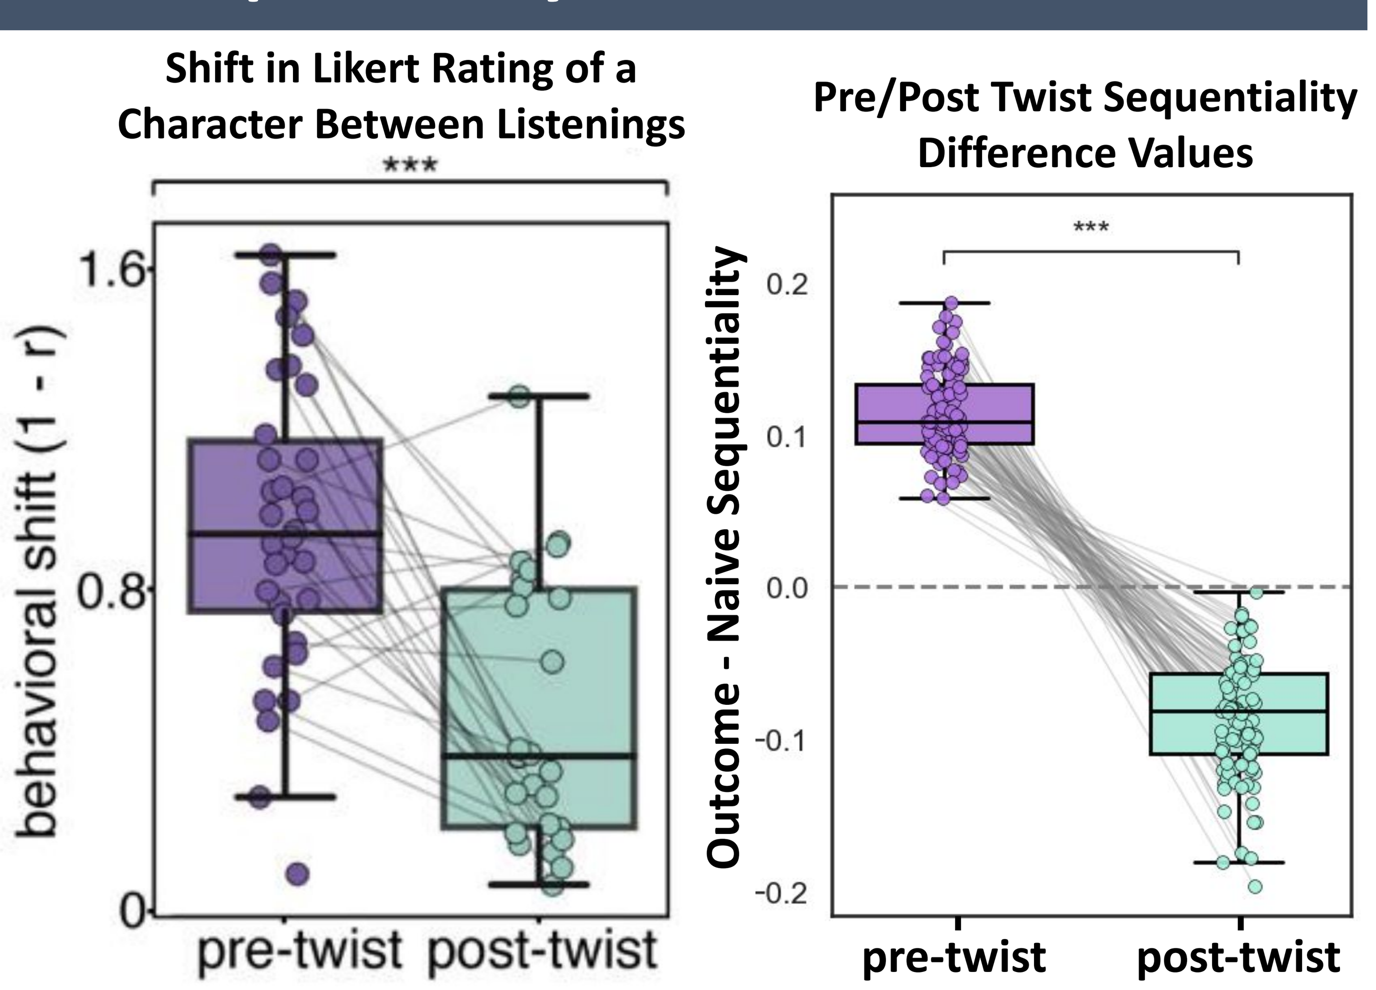


Movie
Annotations
Sequentiality (c)
Event 1: A woman stands alone in a room, opens a box of pizza, and takes a slice.
Event 2: She finds a record player, places a record on it, and lowers the needle.
Event 3: The woman enjoys cheerful music as she eats the slice of pizza.
c(s1) = 1.8
c(s2) = 0.62
c(s3) = 1.7
Mean Importance
Causal Page Rank
Semantic Page Rank
[5, 3, 6, 4] = 4.5
[2, 1, 3, 2] = 2
[2, 4, 3, 5] = 3.5
Σ cos(s1,sn) = 0.05
Σ cos(s2,sn) = 0.02
Σ cos(s3,sn) = 0.04
Σ cos(s1,sn) = 0.04
Σ cos(s1,sn) = 0.04
Σ cos(s1,sn) = 0.02

Research Question

How does our metric of contextual information gain relate to human judgements of narrative structure while viewing naturalistic narrative stimuli?

Sequentiality in Audio Narratives



Shift in Likert Rating of a Character Between Listenings

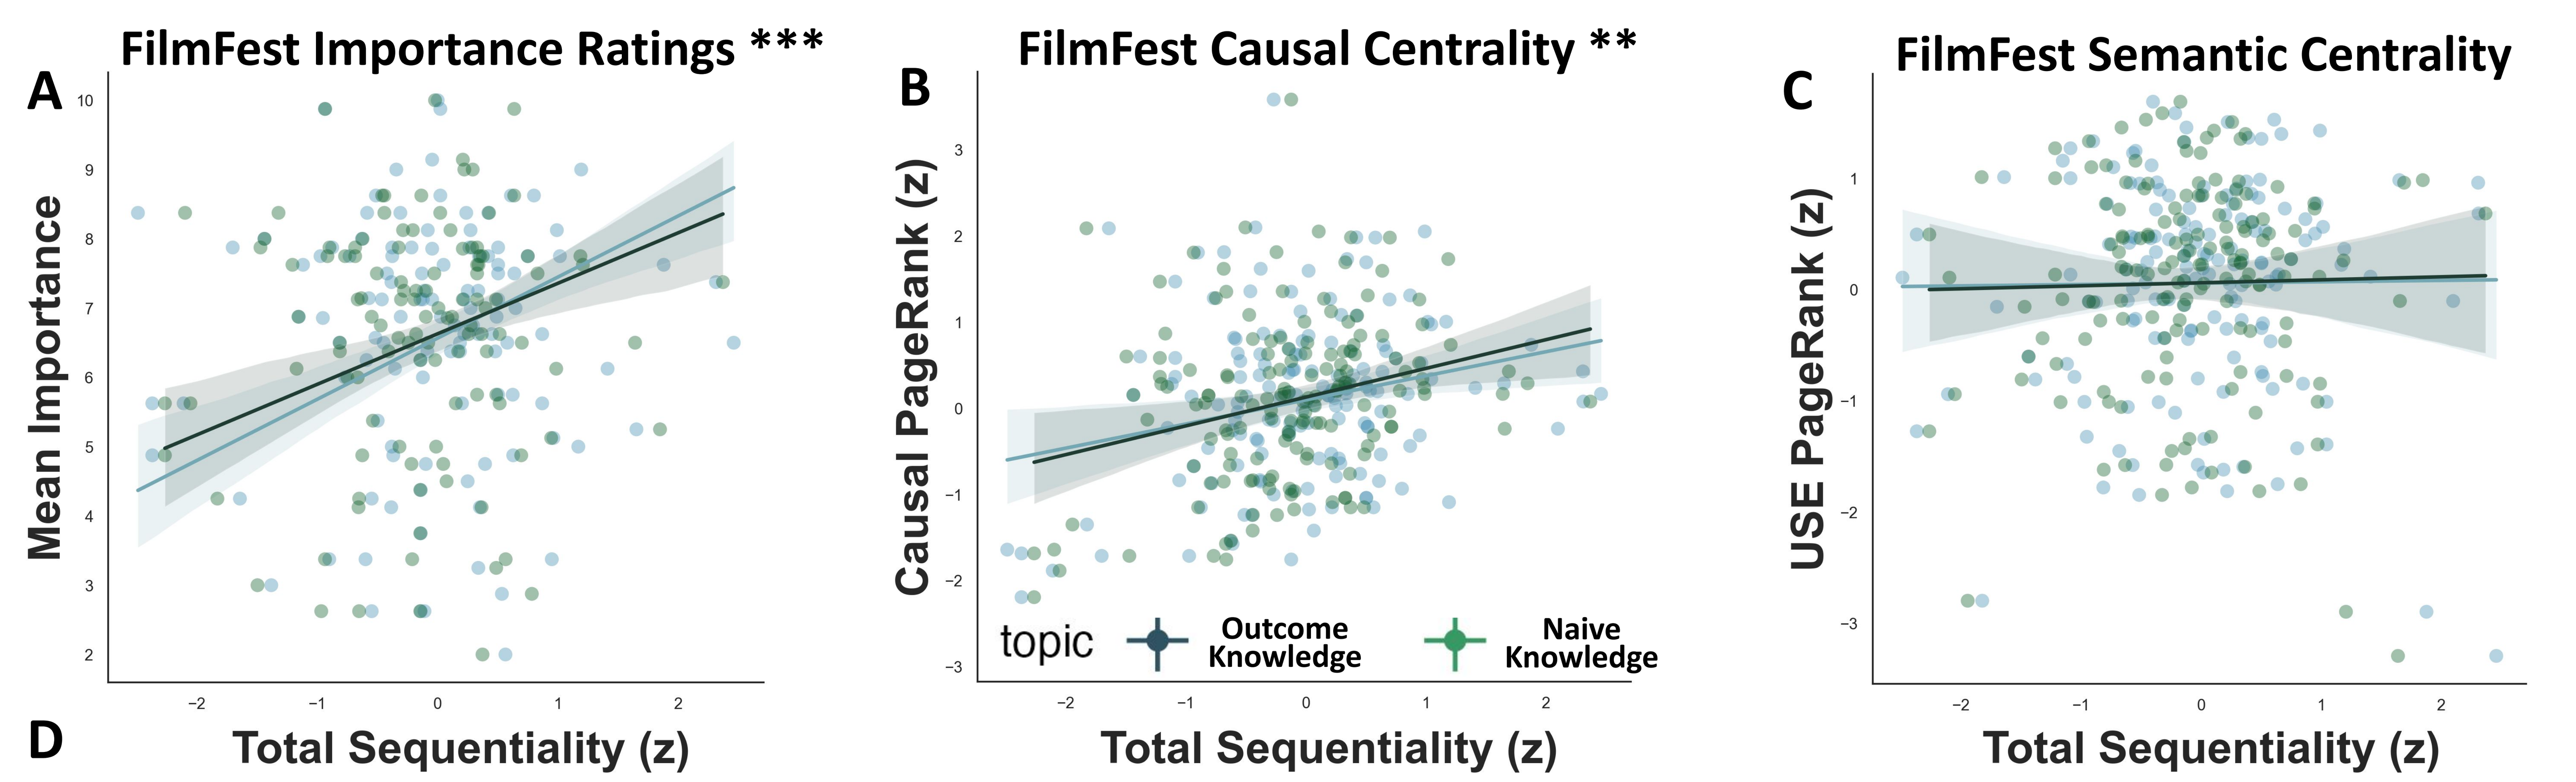
behavioral shift (1 - r)
1.6
0.8
0
pre-twist
post-twist
Pre/Post Twist Sequentiality Difference Values

Outcome - Naive Sequentiality
0.2
0.1
0.0
-0.1
-0.2
pre-twist
post-twist

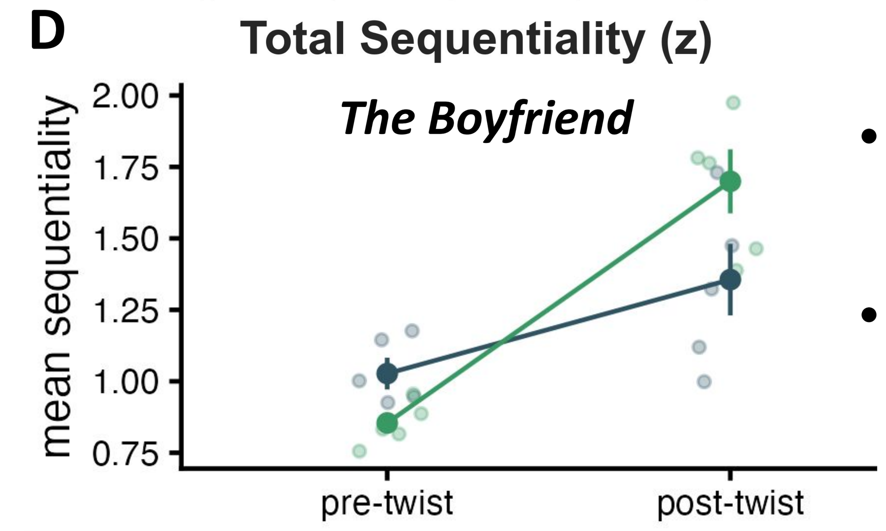
(Sava-Segal et al., 2025)

- Participants' evaluation of a character changed significantly more for pre-twist events during second listening (Sava-Segal et al., 2025)
- Difference between outcome knowledge and naive Sequentiality is *also* significantly higher for pre-twist sentences (t=41.61, p<.001)

Sequentiality Tracks Importance and Causality but not Semantic Centrality



A
FilmFest Importance Ratings ***
Mean Importance
Total Sequentiality (z)
B
FilmFest Causal Centrality **
Causal PageRank (z)
topic
Outcome Knowledge
Naive Knowledge
Total Sequentiality (z)
C
FilmFest Semantic Centrality
USE PageRank (z)
Total Sequentiality (z)



D
The Boyfriend
mean sequentiality
2.00
1.75
1.50
1.25
1.00
0.75
pre-twist
post-twist

- Sequentiality is positively correlated with mean importance values (r=.36, p=.001) **(A)** and Causal PageRank (r=.24, p=.003) **(B)**, and shows no correlation with USE PageRank (r=.01, p=.894) (all for naive, but significance consistent across topics) **(C)**
- Sequentiality values were calculated for all sentences in a story with a twist ending (*The Boyfriend)*. Naive Sequentiality values were higher after the twist, consistent with our information gain hypothesis (OLS regression interaction; β=0.5165, t=4.306, p=.001) **(D)**

Conclusion

Sequentiality is sensitive to event importance, causal structure, and event segmentation across various narratives. Furthermore, the presence of outcome knowledge changes the metric's expectation during twists, similar to human shifts in narrative ratings.

Contact Us!

Contact Info | Preprint | Sequentiality

Model Reformulation

"How well does the story flow?"

$$c(s_i, h) = -\frac{1}{|s_i|}(\log p(s_i|\tau) - \log p(s_i|\tau, h))$$

(Sap et al., 2022)

Combine log terms

$$c(s_i, h) = -\frac{1}{|s_i|}\log\frac{p(s_i|\tau)}{p(s_i|\tau, h)}$$

Expand with Bayes Rule

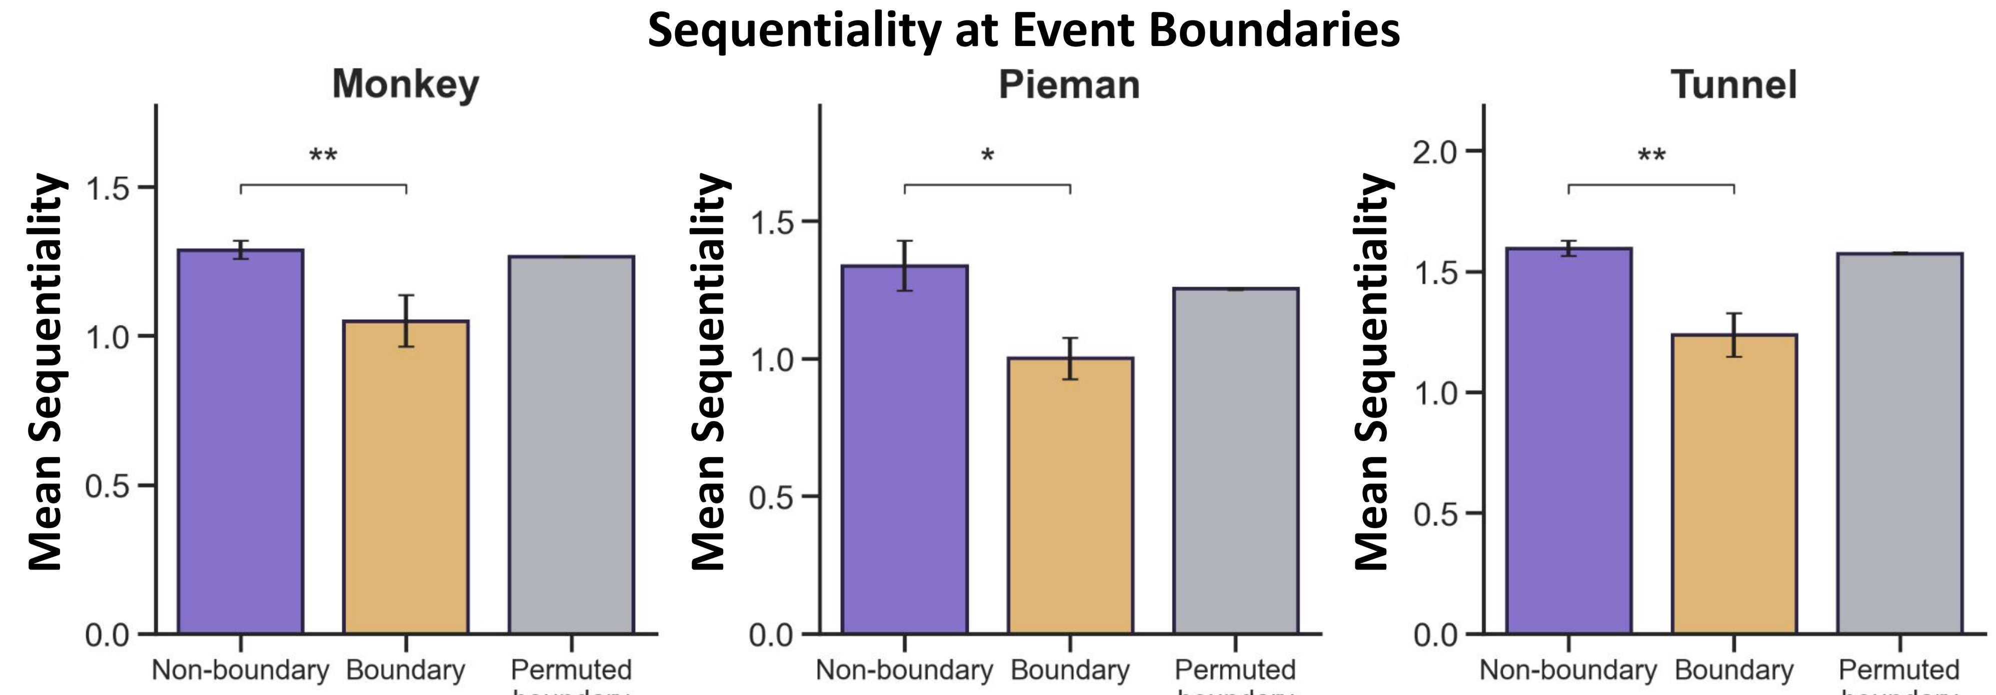
$$c(s_i, h) = -\frac{1}{|s_i|}\log\left(\frac{\frac{p(\tau|s_i)p(s_i)}{p(\tau)}}{\frac{p(\tau,h|s_i)p(s_i)}{p(\tau,h)}}\right)$$

Simplify

$$c(s_i, h) = -\frac{1}{|s_i|}\log\left(\frac{p(h|\tau)}{p(h|\tau, s_i)}\right)$$

"How much does s_i tell you about h beyond what the τ tells you?"

Sequentiality is Sensitive to Event Boundaries



Sequentiality at Event Boundaries
Monkey
Pieman
Tunnel
Mean Sequentiality
Non-boundary
Boundary
Permuted boundary
**
*
**

- Ran Sequentiality on *Pieman, Monkey in the Middle, and Tunnel Under the World* and binned Sequentiality values according to whether the sentence was consistently rated as an event boundary by human subjects (Michelmann et al., 2025)
- Sequentiality is much lower at boundaries than non-boundaries across stories (t=-2.453±0.214)

Future Directions

- Investigate if value processing regions are associated with contextual information gain
- Apply to long timescale narratives where credit assignment is complex and nonmonotonic
- Use Sequentiality to guide generation of narrative stimuli that can test effects of causal structure on memory and choice

References

Antony, J., Lozano, A., Dhoat, P., Chen, J., & Bennion, K. (2024). Causal and chronological relationships predict memory organization for nonlinear narratives. J. Cogn Neurosci., 36(11), 2368–2385. Chen, J., & Bornstein, A. M. (2024). The causal structure and computational value of narratives. Trends in Cognitive Sciences, 28(8), 769–781. DuBrow, S., Rouhani, N., Niv, Y., & Norman, K. A. (2017). Does mental context drift or shift? Curr Opin Behav Sci, 141–146. Goeschel, A., Palsule, R., Chen, J., Bornstein, A.M. Khoudary, A. (2026) Quantifying information gain in narrative stimuli using open-source large language models:preprint. Lee, H., & Chen, J. (2022). Predicting memory from the network structure of naturalistic events. Nature Communications, 13, 4235. Nguyen, T. T., Etzel, J. A., Bezdek, M. A., & Zacks, J. M. (2026). Multiple event segmentation mechanisms in the human brain. eLife, 13. Panela, R. A., Barnett, A. J., Barense, M. D., & Herrmann, B. (2025). Event segmentation applications in large language model enabled automated recall assessments. Commun. Psychol., 3(1), 184. Raccah, O., Chen, P., Gureckis, T. M., Poeppel, D., & Vo, V. A. (2024). The "Naturalistic Free Recall" dataset: Four stories, hundreds of participants, and high-fidelity transcriptions. Scientific Data, 11(1317). Sap, M., Jafarpour, A., Choi, Y., Smith, N. A., & Horvitz, E. (2022). Quantifying the narrative flow of imagined versus autobiographical stories. Proceedings of the National Academy of Sciences of the United States of America, 119(45). Sava-Segal, C., Grall, C., & Finn, E. S. (2026). Narrative "twist" shifts within-individual neural representations of dissociable story features. Proceedings of the National Academy of Sciences, 123(11). Yoo, J. (2025). Contingency as a fundamental principle for learning cognitive graphs. ProQuest.

Acknowledgements

Thank you to the Chen Lab at JHU and the UCI Computational Cognitive Neuroscience Lab for their assistance. Additionally, special thanks to Dr. Maarten Sap and Dr. Anna Jafarpour for introducing this metric and assisting with the model reformulation.