Context modulates brain state dynamics and behavioral responses during narrative comprehension

Yibei Chen[1*], Zaid Zada[2], Samuel A. Nastase[2], F. Gregory Ashby[3], Satrajit S. Ghosh[1]

[1]McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, MA, 02139, USA

[2] Princeton Neuroscience Institute, Princeton University, Princeton, NJ, 08544, USA

[3] Department of Psychological & Brain Sciences, University of California, Santa Barbara, Santa Barbara, CA, 93106, USA

*For correspondence: yibei@mit.edu

Abstract

Narrative comprehension is inherently context-sensitive, yet the brain and cognitive mechanisms by which brief contextual priming shapes story interpretation remain unclear. Using hidden Markov modeling (HMM) of fMRI data, we identified dynamic brain states as participants listened to an ambiguous spoken story under two distinct narrative contexts (affair vs. paranoia). We uncovered both context-invariant states—engaging auditory, language, and default mode networks—and context-specific states characterized by differential recruitment of control, salience, and visual networks. Narrative context selectively modulated the influence of character speech and linguistic features on brain state expression, with the central character's speech enhancing activation in shared states but suppressing activation in context-specific ones. Independent behavioral analyses revealed parallel context-dependent effects, with character-driven features exerting strong, selectively modulated influences on participants' judgments of narrative evidence. These findings demonstrate that brief narrative priming actively reshapes brain state dynamics and feature sensitivity during story comprehension, revealing how context guides moment-by-moment interpretive processing in naturalistic settings.

Keywords: narrative comprehension, context modulation, brain state dynamics, hidden Markov models, Bayesian mixed-effects modeling, naturalistic paradigm, fMRI

Introduction

Narrative comprehension involves complex interactions among prior knowledge, immediate contextual expectations, and the content (Botch & Finn, 2024; Mar & Oatley, 2008;

Nastase et al., 2021; Willems et al., 2020). Recent neuroimaging research demonstrates significant variability in how individuals process identical narrative stimuli, primarily driven by stable personal traits such as empathy, political beliefs, and cognitive abilities. This variability results in distinct patterns of brain activity and synchronization (Coderre & Cohn, 2023; de Bruin et al., 2023; Johns et al., 2018; Nijhof & Willems, 2015). In addition to these stable individual differences, transient manipulations of expectations or interpretation profoundly impact how narrative content is processed, highlighting the brain's sensitivity to context (Yeshurun et al., 2017).

Research in psychology and linguistics highlights the crucial role of context in understanding narratives. Contextual cues, such as background knowledge or primed information, help people construct coherent mental representations, improve comprehension, and enhance recall of narrative details (Bransford & Johnson, 1972; van Kesteren et al., 2013; Zwaan & Radvansky, 1998). Neuroimaging studies further show that naturalistic narratives (e.g., audiostories or movies), which unfold over time, recruit default mode network regions that are typically less engaged by highly controlled lab stimuli that are shorter and decontextualized (Baldassano et al., 2017; Ben-Yakov et al., 2012; Geerligs et al., 2022; Lerner et al., 2011; Yeshurun et al., 2021). These longer narratives allow context to accumulate gradually, shaping interpretation as the story progresses and enabling the construction of shared mental models across listeners.

When individuals receive similar contextual information, whether through common prior knowledge or priming, their cognitive and emotional responses tend to align, producing synchronized activity across brains (de Bruin et al., 2023; Lahnakoski et al., 2014; Nguyen et al., 2019). Conversely, when people are primed differently or bring distinct prior experiences to the same narrative, they may interpret it in diverging ways, resulting in idiosyncratic patterns of brain activity (Jacoby & Fedorenko, 2020; Yeshurun et al., 2017). These interpretive processes rely on integrating incoming information with existing mental models and are thought to emerge from coordinated activity across multiple large-scale brain networks, rather than being localized to any single region (Barrett, 2022; McIntosh, 2004; Song et al., 2023).

Narrative comprehension relies on relatively stable cognitive frameworks, such as schemas or prior knowledge about story structure (Mar, 2011). Still, it can also be flexibly shaped by short-term contextual manipulations, such as priming just before the narrative begins.

These immediate contextual cues influence how listeners engage with incoming narrative content by biasing expectations and interpretive stance from the outset. The mechanisms by which such external contextual priming interacts with the brain's ongoing maintenance of context remain poorly understood. In this study, we examine how the integration of narrative input with initial contextual priming is reflected in dynamic patterns of brain activity, using the concept of "brain states." Brain states refer to recurring patterns of coordinated activity across distributed brain regions (Liu et al., 2025; Song et al., 2021), analogous to distinct musical motifs formed by different instruments in an orchestra. By identifying and characterizing these patterns, we can assess the temporal dynamics of brain activity during narrative processing and how they differ across primed context conditions (Shine et al., 2019; Vidaurre et al., 2017).

Similarly, a critical gap remains in understanding the brain mechanisms by which contextual priming interacts with specific narrative features, such as character identity and other linguistic structures. Character identity is conveyed and reinforced during narrative processing through character speech, particularly when direct speech is attributed to specific characters. These attributions fundamentally shape comprehension by guiding attention, emotional engagement, and social inference (Gerrig, 1993; Jacoby & Fedorenko, 2020; Mar & Oatley, 2008). Psycholinguistic evidence consistently underscores the central role of character speech in maintaining narrative coherence, supporting mental simulation, and enabling theory-of-mind reasoning (Nieuwland & Van Berkum, 2006; Zwaan & Radvansky, 1998). Thus, character speech serves as a theoretically meaningful and empirically tractable feature for investigating how contextual priming influences narrative processing.

To investigate how narrative context shapes brain state dynamics during story comprehension, we used a naturalistic fMRI paradigm in which two participant groups listened to the same story but were primed differently beforehand (Yeshurun et al., 2017). We applied hidden Markov models (HMMs) to identify recurrent brain states, defined as temporally evolving patterns of network-level activity, across the full duration of story listening (Quinn et al., 2018; Shine et al., 2019; Taghia et al., 2018; Vidaurre et al., 2017). We hypothesized that narrative comprehension involves both stable primary states that track core story structure and flexible secondary states that vary with contextual integration demands, particularly in networks involved in attention, visual processing, and cognitive control (Baldassano et al., 2017; Hasson et al., 2015). Building on prior work showing brain sensitivity to character-level features

(Alderson-Day et al., 2020; Jacoby & Fedorenko, 2020; Yarkoni et al., 2008), we further predicted that brain state dynamics would differ based on speaker identity.

We added a complementary behavioral experiment to better understand how primed context influences moment-to-moment interpretation. Our goal was to capture when listeners subjectively recognized elements of the story as aligning with their assigned context. Two separate groups of participants received the same context instructions and listened to the same story as those in the fMRI study. They were asked to press a key whenever they perceived information consistent with their contextual framing. These responses provide a time-resolved behavioral index of interpretive alignment, offering an external marker of how context interacts with narrative features over time.

Methods

fMRI dataset

We utilized the “prettymouth” dataset (Figure 1) (Yeshurun et al., 2017), which includes 40 participants drawn from the Narratives data collection (Nastase et al., 2021). Participants were divided into two groups (initially N = 20 per group), with both groups exposed to an adapted version of J. D. Salinger's short story, "Pretty Mouth and Green My Eyes." The adapted version was shorter than the original and included several sentences not present in the original text. A professional actor provided the narration, resulting in a recording of 11 minutes and 32 seconds. Functional MRI data were acquired with a repetition time (TR) of 1.5 seconds. The story was preceded by 18 seconds of neutral music and 3 seconds of silence, followed by an additional 15 seconds. These segments of music and silence were excluded from all analyses.

The narrative describes a phone conversation between two friends, Arthur and Lee. Arthur, who has just returned home from a party after losing track of his wife Joanie, calls Lee to express his concerns about her whereabouts. Lee is at home with a woman beside him, whose identity remains ambiguous: she may or may not be Joanie. Before listening to the story, each participant group received one of two different contextual prompts: one group was informed that Arthur was paranoid and his suspicions were unfounded (“paranoia” context), while the other group was told that the woman was indeed Joanie, Arthur’s wife, and that Lee and Joanie had been involved in an ongoing affair for over a year (“affair” context). Yeshurun et al. (2017) and Nastase et al. (2021) describe the experimental paradigm and fMRI data acquisition parameters.

fMRI data processing

fMRI data preprocessing was conducted using fMRIPrep (version 24.0.1) (Esteban et al., 2019) via the BIDS App Bootstrap (Zhao et al., 2024), with detailed processing steps provided in the supplementary material. Custom post-processing steps optimized for narrative listening analyses were applied after initial preprocessing. Specifically, we implemented spatial smoothing (6 mm full-width half-maximum) to balance noise reduction and preservation of spatial activation patterns, performed detrending to mitigate scanner drift, standardized (z-scored) the time-series signals across time points within each voxel within each subject, and regressed out nuisance signals related to head motion and physiological noise using motion parameters and anatomical CompCor regressors. All post-processing steps were carried out using Nilearn; more details are provided in the supplementary material and our GitHub repository (see the Code Availability section). Post-processed data were further extracted using the Schaefer et al. (2018) parcellation (1000 parcels) with 17 networks (Kong et al., 2021). Following recommendations by Nastase et al. (2021), two participants were excluded from further analysis due to data quality concerns, resulting in a final sample size of N = 19 per group.

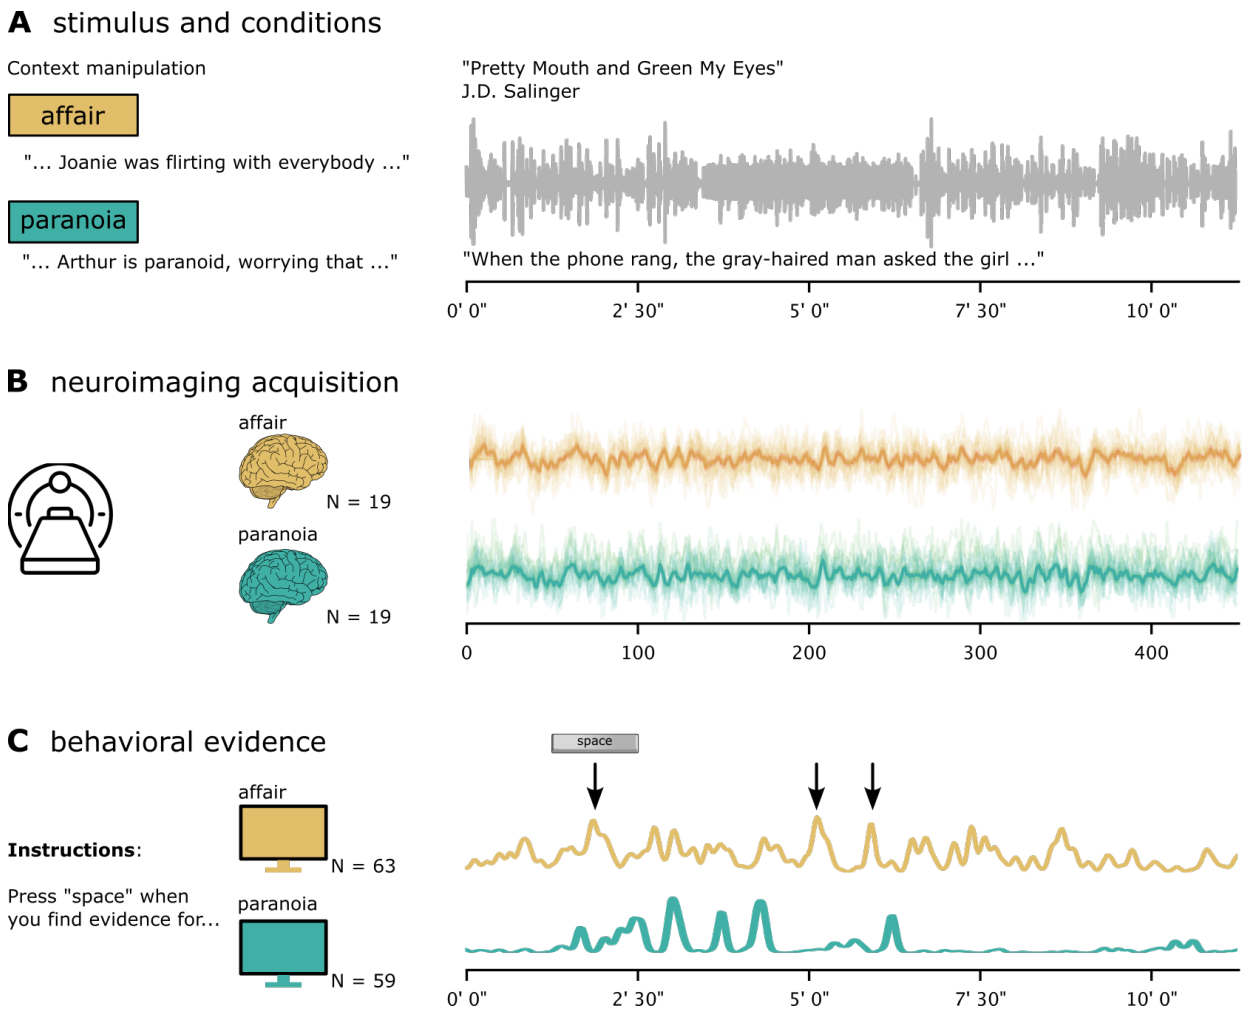

Figure 1. Experimental design, neuroimaging acquisition, and behavioral evidence. (**A**) Participants were randomly assigned to one of two context conditions—affair (gold) or paranoia (teal)—and read a brief prompt before listening to an 11-minute spoken story (Pretty Mouth and Green My Eyes by J.D. Salinger). This context manipulation was consistent across both neuroimaging and behavioral experiments. (**B**) The fMRI study included data from 19 participants in each context group in the final analysis. The schematic plot illustrates average fMRI time courses for each group (affair in gold, paranoia in teal), with the x-axis representing time in TRs. (**C**) In the behavioral study, we recruited two new groups of participants and asked them to press the spacebar whenever they perceived evidence supporting one interpretation or the other (affair N = 63, paranoia N = 59). The line plots depict the average frequency of button presses over time within each group (i.e., agreement across all participants), with peaks corresponding to moments of perceived narrative support for each interpretation.

Behavioral data

We collected an additional behavioral dataset under two different tasks to assess the evidence in the stimulus supporting each narrative context over time. Behavioral data were collected from 128 participants recruited via Prolific (www.prolific.com). Participants were classified into two experimental groups: an affair group (n = 63) and a paranoia group (n = 59). Demographic details indicated the sample comprised 62 males, 57 females, and three individuals identifying as non-binary or third gender. The age distribution included participants aged 18–24 years (n = 18), 25–34 years (n = 46), 35–44 years (n = 29), 45–54 years (n = 12), 55–64 years (n = 14), and 65 years or older (n = 3). These participants are a separate sample from those included in the fMRI experiment.

Data were collected through an online experiment developed using PsychoJS scripts derived from the PsychoPy builder (PsychoPy3, version 2023.2.0), hosted on Pavlovia (https://pavlovia.org/). Participants initially provided informed consent via Qualtrics (https://www.qualtrics.com/) before being randomly assigned to one of two context conditions (affair versus paranoia). Participants in each group received the same prompts presented to the fMRI participants prior to listening to the auditory story. Participants were asked to identify moments in the narrative where they perceived evidence for their assigned interpretation (Lee and Joanie are having an affair, or Arthur is being paranoid) by pressing the spacebar on their keyboards. Immediate visual feedback was provided, indicated by a brief green dot appearing at the center of the screen, confirming each response. After completing the task, participants were redirected to Qualtrics to complete a post-experiment questionnaire. Data from four participants were excluded from subsequent analyses due to incomplete records, resulting in a final dataset of 122 participants (Figure 1C). More detailed instructions can be found in the supplementary material.

This study was approved by the Princeton University Institutional Review Board (IRB 12201). In accordance with institutional ethical guidelines, all participants provided informed consent electronically before participation. Participants received monetary compensation consistent with university policy. All data were anonymized to ensure participant confidentiality.

Hidden Markov model (HMM) analysis

To characterize the temporal dynamics of brain states during story listening, we employed hidden Markov models (HMMs, Figure 2A), which identify recurring patterns of brain network activity and their transitions over time (Baldassano et al., 2018; Meer et al., 2020; Vidaurre et al., 2017; Yang et al., 2023). HMMs explicitly model temporal dependencies and sequential state transitions, aligning closely with our objective of understanding how prior contextual information modulates the temporal evolution of brain states.

To account for hemodynamic delay, the BOLD signal was shifted backward by three TRs (~4.5 s) relative to the timing of the stimulus features (Yeshurun et al., 2017). Non-story segments (background music/silence, 24 TRs at scan onset and offset) were excluded, yielding 451 TRs for analysis. Time series were extracted from 17 functionally defined networks by first averaging voxel-wise signals within each parcel, then averaging across all parcels assigned to the same network. Each participant's data were z-scored to normalize signal amplitudes.

We implemented Gaussian observation HMMs using the *hmmlearn* Python package, modeling brain states as multivariate Gaussian distributions with state-specific means and covariances. Transition probabilities were initialized to favor an expected dwell time of approximately 7 seconds (~4–5 TRs). This provided a weak prior consistent with evidence that large-scale brain states during naturalistic cognition typically persist on the order of several seconds, with higher-order regions such as the default mode network exhibiting longer durations and sensory regions (e.g., auditory cortex) switching more rapidly (Vidaurre et al., 2017; Baldassano et al., 2017). The prior served to prevent implausibly fast state-switching at initialization, but transition probabilities were re-estimated during expectation–maximization, so final dwell times were determined by the data rather than fixed to 7 seconds.

To improve robustness against local optima during model fitting, we performed five independent initialization attempts per model (i.e., with the same number of states but different random seeds). For each attempt, state means were initialized by random draws from a standard normal distribution ($\mathcal{N}(0,1)$), ensuring broad coverage of the feature space. Covariance matrices were initialized as identity matrices with a small diagonal ridge term ($\epsilon = 10^{-6}$) to maintain positive definiteness and numerical stability during optimization. Start probabilities were initialized uniformly across states, and transition priors followed the weak dwell-time prior described above. Across the five attempts, random seeds influenced both the initialization of

means and the stochastic path of the EM algorithm, while covariances, start probabilities, and transition priors were held constant. The solution with the highest log-likelihood was selected for downstream analysis.

Model generalizability and stability were assessed through leave-one-subject-out cross-validation (LOOCV).For each fold, a group-level model was trained on all but one subject and evaluated on the held-out subject's data. We quantified model fit using the average cross-validated log-likelihood across folds. State reliability across folds was evaluated via spatial correlations, with optimal state matching determined using the Hungarian algorithm(Kuhn, 1955). Pattern similarity scores were averaged across all fold pairs to yield a summary measure of cross-validated pattern reliability. We also provided the model selection result from LOOCV in **Supplementary Figure 1**.

HMM analyses were conducted on two experimental groups—the affair context group (n=19) and paranoia context group (n=19)—and two constructed groups—a combined group comprising all participants (n=38) and a balanced group (n=19), to match the contextual groups, consisting of random subset of participants (9 from the affair context and 10 from the paranoia context). The balanced group was created to preserve the contextual heterogeneity of the combined group while matching the sample size of the individual context groups. This approach enabled both the investigation of context-specific brain state dynamics and generalizable patterns across contexts. Brain states were characterized by spatial patterns, temporal sequence (i.e., occurrence), and inter-subject consistency. All analyses were implemented using Python (hmmlearn, NumPy, SciPy).

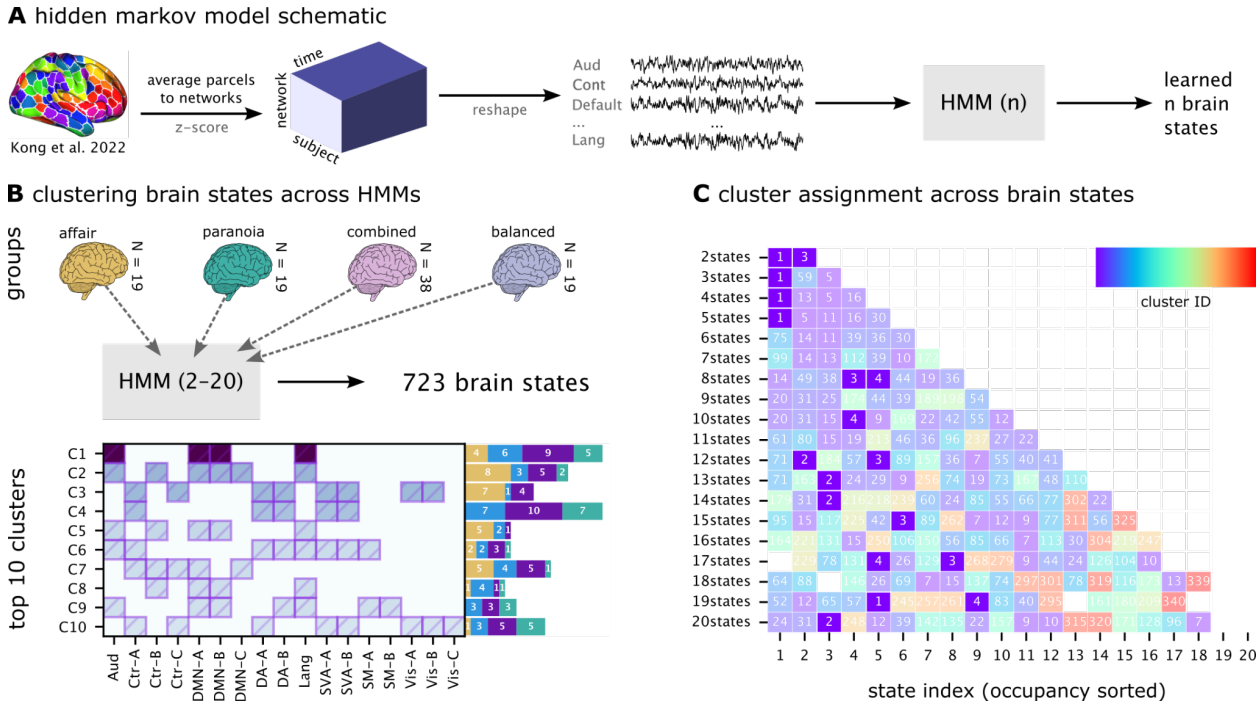

Figure 2. Hidden Markov models (HMMs) and clustering of brain states. (**A**) Analysis pipeline. Parcel-level fMRI time series were averaged within networks, z-scored across time, and concatenated across participants. HMMs with 2–20 states were estimated separately for four groups: *affair*, *paranoia*, *combined*, and *balanced*. (**B**) Cluster profiles and provenance. Brain states that passed quality filters (activation > 0.1, bootstrap CI width < 0.3, split-half reliability > 0.5) were pooled across models and clustered by spatial similarity (Jaccard distance, hierarchical clustering). The left panel shows the average activation profile of each resulting cluster across canonical functional networks, with shading intensity indicating the mean activation level. The bar plots to the right show the number of states from each group (affair, paranoia, combined, balanced) contributing to each cluster, illustrating cluster provenance. (**C**) Cluster assignments across model granularities. Results are shown for the *combined*-group HMMs. Each row corresponds to one HMM solution (2–20 states), and each column to a state within that model, ordered by state occupancy. Colors and overlaid numbers indicate the cluster assignment (1–4) of each state. Missing cells indicate states that did not pass the filtering criteria and were therefore excluded from clustering. This panel demonstrates how states from different HMM solutions converge onto a consistent set of four clusters.

State pattern similarity analysis

Traditional methods for analyzing hidden Markov models usually depend on model selection criteria to pinpoint a single optimal model, often overlooking valuable insights from alternative solutions (Quinn et al., 2018; Vidaurre, Hunt, et al., 2018). To overcome this limitation, we created a pattern similarity analysis framework that utilizes information from various model parameterizations and experimental conditions. Our approach builds on the neurobiological observation that increasing the number of states in an HMM often results in meaningful subdivisions of broader brain state processes rather than entirely spurious patterns (Baker et al., 2014). For example, a language processing state in a simpler model might subdivide into different states for different semantic domains in more complex models, representing valid phenomena at different levels of granularity.

Reliable state patterns from all HMM solutions across the experimental (affair, paranoia) and constructed groups (combined, balanced) were extracted using stringent criteria designed to filter out unstable or noisy states. First, we required a minimum activation threshold (>0.1), meaning that at least one network's mean parameter in the HMM had to deviate from baseline by more than 0.1 (in standardized units of the z-scored input time series). This ensured that states reflected meaningful network engagement rather than near-zero fluctuations. Second, we assessed the stability of each state's mean activation pattern using bootstrap resampling: for each network, we generated 95% confidence intervals around its mean activation value and retained only states with narrow intervals (width < 0.3), indicating low estimation uncertainty. Third, we evaluated within-state reliability using a split-half procedure: timepoints assigned to each state were divided into two halves, mean activation patterns were estimated separately, and their correlation was computed; only states with correlations >0.5 were retained. For each retained state, we recorded fractional occupancy to standardize comparisons across models, along with full provenance information (group, model specification, original state index, normalized index).

We then clustered the state mean activation patterns across groups (Figure 2B), focusing on the average co-activation profile of the 17 networks for each state. Pairwise similarity was computed using Jaccard distance, which quantifies dissimilarity based on the proportion of non-overlapping active networks. This metric emphasizes the spatial layout of co-active brain networks while reducing sensitivity to overall activation magnitude. All states were compared pairwise, and agglomerative hierarchical clustering with average linkage was applied to group

them. We evaluated clustering solutions across merging similarity thresholds from 0.6 to 0.9 and found that the top four clusters were highly stable across this range (**Supplementary Figure 3**). For subsequent analyses, we used the 0.8 merging threshold.

As part of the cross-group clustering procedure, clusters were reordered by total fractional occupancy, defined as the sum of the occupancies of all states assigned to each cluster across models. This ensured that cluster IDs reflected the most frequently expressed patterns rather than simply the largest number of constituent states. The occupancy-based reordering occurred immediately after hierarchical clustering, integrated within the clustering step rather than as a separate stage. Representative brain states for each cluster were then calculated by averaging the patterns of states within the cluster and identifying features consistently activated in at least half of the contributing states.

Finally, clusters were labeled as either context-general or context-specific based on their provenance. Clusters containing states from both affair and paranoia models (and also appearing in the combined/balanced models) were considered context-general, whereas clusters dominated by states from only one context group were considered context-specific. To evaluate the similarity structure among brain state clusters, we computed Spearman rank correlations of their activation patterns across the 17 networks. We conducted the analysis using both consensus patterns (averaged across all member states) and representative patterns (from the combined-group models), which yielded convergent results. Here we report the correlations based on representative patterns; full results from both approaches are provided in **Supplementary Figure 8**.

Story feature annotation

To examine how narrative content influenced brain state dynamics, we annotated the stimulus with key linguistic and narrative features at the temporal resolution of the fMRI data (one annotation per TR).

Character and interaction features: As the narrative was delivered by a single narrator but featured multiple characters, we identified character-specific speech and interactions per TR. The annotations included: (1) Arthur, Lee, and Girl speaking: Identify the intended speaking character at each time point. (2) Lee and the girl together: Identify when Lee and the girl appeared concurrently, regardless of dialogue.

Linguistic features: Story were tagged for grammatical parts of speech, including verbs, nouns, adjectives, and adverbs, indicating their presence in each TR.

Thematically relevant combined features: We further derived composite features to reflect interactions between character presence and linguistic structure: (1) Lee-Girl Verb (Lee & Girl Together × Verb Presence): Captured shared actions or relational dynamics relevant to the affair group's expected sensitivity to relational events. (2) Arthur Adjective (Arthur Speaking × Adjective Presence): Highlighted descriptive attributes linked to Arthur, informed by findings that heightened attention to character traits is characteristic of paranoid cognition (M. J. Green & Phillips, 2004).

These structured annotations enabled the systematic evaluation of how different narrative elements influenced cognitive engagement, providing an essential foundation to investigate the hypothesized cognitive biases associated with each group.

Bayesian generalized linear mixed models

To investigate the temporal dynamics of brain state patterns and corresponding behavioral responses, and to clarify how contextual information modulates the impact of narrative content features, we implemented Bayesian generalized linear mixed models (GLMMs). Separate GLMM analyses with identical structures were applied to characterize brain-context-content and behavior-context-content relationships, providing consistent modeling frameworks for brain and behavioral dynamics. The schematic overview of the full pipeline, from HMM to GLMM, is in **Supplementary Figure 2**.

GLMM for brain state and content analysis

While the clustering analysis identified spatial configurations of brain states, the temporal dynamics necessitated identifying representative state occurrences. A representative brain state was chosen for each cluster's first occurrence within the combined group HMMs, as these models included all participants. Subsequently, we extracted each participant's state sequence (on/off) data corresponding to these representative states at each time point. We fit a logistic GLMM separately for each identified cluster with the following structure:

$$\mathrm{logit}(P(\mathrm{State}_{it}=1)) = \beta_0 + \beta_g \cdot \mathrm{Group}_i + \sum_{j=1}^{J} \beta_j \cdot \mathrm{Feature}_{jt} + \sum_{j=1}^{J} \beta_{gj} \cdot \mathrm{Group}_i \cdot \mathrm{Feature}_{jt} + \sum_{k=1}^{2} \gamma_k \cdot \mathrm{State}_{i,t-k} + u_i$$

Where State_{it} is a binary variable indicating whether the target brain state was active (1) or inactive (0) for subject i at timepoint t. $\mathrm{Feature}_{jt}$ represents narrative annotations (e.g.,

character presence, linguistic elements). The interaction terms $\mathrm{Group}_i \cdot \mathrm{Feature}_{jt}$ assess whether content features affect brain state dynamics differently between groups. Autoregressive terms $\mathrm{State}_{i,t-k}$ account for temporal dependencies in state occupancy, and u_i represents subject-specific random intercepts that capture individual variability in state prevalence.

Model parameters were estimated using maximum a posteriori (MAP) estimation, applying deviation coding for group identity (+1 for affair, -1 for paranoia) and incorporating default Bayesian priors: normal priors with a mean of 0 for fixed effects and inverse gamma priors for random effects variance components. These priors provide implicit regularization, which is advantageous given our moderate sample size and binary outcomes. We calculated posterior probabilities instead of frequentist p-values for inference, quantifying evidence for effects as the probability mass supporting a specific direction of influence. This Bayesian approach allows for a more intuitive interpretation of uncertainty in our parameter estimates.

To address multiple comparisons, we implemented a Bayesian False Discovery Rate (FDR) procedure that controls the expected proportion of false discoveries among claimed discoveries. Features were considered to have credible effects when their FDR-adjusted posterior probabilities exceeded 0.95.

Coefficient estimates were converted from log-odds to odds ratios (OR) to enhance interpretability, indicating how narrative features influenced primary brain state activation odds. Group-specific effects were calculated to clarify how content features differentially affected brain state dynamics in each context condition. All analyses were performed using custom Python with the statsmodels package.

GLMM for behavioral response and content analysis

We applied a generalized linear mixed model (GLMM), analogous to those used in the brain state analyses, to examine the relationship between narrative content features and behavioral responses in a separate participant sample. The dependent variable was a binary indicator reflecting whether a button press occurred at each fMRI time point (TR), signaling that the participant perceived evidence in the narrative consistent with their assigned contextual prompt. Originally recorded continuously (seconds), behavioral responses were aligned to the nearest TR to ensure temporal correspondence with stimulus features and brain-state estimates. If multiple button presses occurred within a single TR for a given participant, they were counted as a single response to avoid overrepresenting clustered inputs.

The behavioral GLMM followed this structure:

$$\text{logit}(P(\text{Response}_{it}=1)) = \beta_0 + \beta_g \cdot \text{Group}_i + \sum_{j=1}^{J} \beta_j \cdot \text{Feature}_{jt} + \sum_{j=1}^{J} \beta_{gj} \cdot \text{Group}_i \cdot \text{Feature}_{jt} + \sum_{k=1}^{2} \gamma_k \cdot \text{Response}_{i,t-k} + u_i$$

Where Response_{it} is a binary variable indicating whether subject i pressed the button at timepoint t. Feature_{jt} represents narrative annotations (e.g., character presence, linguistic elements). Interaction terms $\text{Group}_i \cdot \text{Feature}_{jt}$ assess whether content features differentially affect behavioral responses across groups. Autoregressive terms $\text{Response}_{i,t-k}$ account for temporal dependencies in response patterns, and u_i represents subject-specific random intercepts.

For parameter estimation, we utilized Maximum A Posteriori (MAP) estimation with Bayesian priors to stabilize the estimates, which is particularly important for binary outcomes with temporal dependencies. Similar to the brain-content analyses, we applied deviation coding for group identity (+1 for affair, -1 for paranoia), ensuring that the parameter estimates were balanced around the overall mean. Random intercepts at the subject level accounted for variability in individual response tendencies. Effects were deemed credible when their FDR-adjusted posterior probabilities surpassed 0.95.

Results

Brain state clustering identifies shared and context-specific cortical network patterns

The clustering analysis of brain state patterns across all models identified both context-specific and context-invariant brain states. Models derived from individual experimental conditions (affair or paranoia) produced more context-specific state patterns, as indicated by higher cluster IDs corresponding to smaller clusters, particularly when the number of states increased. Conversely, combined and constructed groups exhibited more generalized state patterns, represented by lower cluster IDs indicating larger clusters (**Supplementary Figures 4-7**). Brain states extracted from the combined group were more consistent than those extracted from the constructed balanced group, likely due to differences in sample sizes. Broad states that dominate at low model complexities (e.g., Cluster 1 in 2–5 state models) can fractionate into multiple, lower-occupancy sub-states at higher model complexities, which explains their apparent “disappearance” in some intermediate state solutions. This likely reflects state fractionation, a fundamental property of hierarchical brain state organization (Li et al., 2023).

The four most prominent clusters (Figure 2C), identified based on the total fractional occupancy of their constituent brain states, displayed distinct spatial configurations. Cluster 1, which had the highest total occupancy (sum=6.288), featured a representative brain state involving the Auditory, DMN-A, DMN-B, and Language networks. Cluster 2, the most frequently occurring cluster, included a representative brain state encompassing Auditory, Control-B, DMN-A, DMN-B, DMN-C, and Language networks.

Clusters 3 and 4 displayed distinct context-specific characteristics. Cluster 3 primarily comprised brain states derived from the affair context models and was largely absent in the paranoia context models. Its representative state pattern included Control-A, Control-C, Dorsal Attention-A, Dorsal Attention-B, Salience/Ventral Attention-A, Salience/Ventral Attention-B, Visual-A, and Visual-B networks. In contrast, Cluster 4 was predominantly composed of brain states from paranoia context models and was absent in the affair context models. The representative pattern of this cluster involved Control-A, Salience/Ventral Attention-A, and Salience/Ventral Attention-B networks. Details of the regions within each network can be found in Table S1.

Spearman's rank correlations revealed positive associations between Clusters 1 and 2 ($\rho = 0.583$) and between Clusters 3 and 4 ($\rho = 0.627$), while Clusters 1/2 versus Clusters 3/4 were negatively correlated. Full correlation results are shown in **Supplementary Figure 8**. Furthermore, **Supplementary Figure 9** illustrates these relationships in low-dimensional space (PCA, MDS, t-SNE), demonstrating clear within-cluster similarity, across-cluster separation, and closer proximity between Clusters 1 and 2.

Contextual modulation of brain state dynamics during story comprehension

Representative brain states (Figure 3) for each identified cluster were selected based on their first occurrence in the combined group HMMs: Cluster 1 (first state from 2-states model), Cluster 2 (fifth state from 7-states model), Cluster 3 (second state from 2-states model), and Cluster 4 (third state from 6-states model). Bayesian GLMMs were then estimated separately for each cluster to determine how narrative features influenced brain state dynamics and whether these effects were modulated by context.

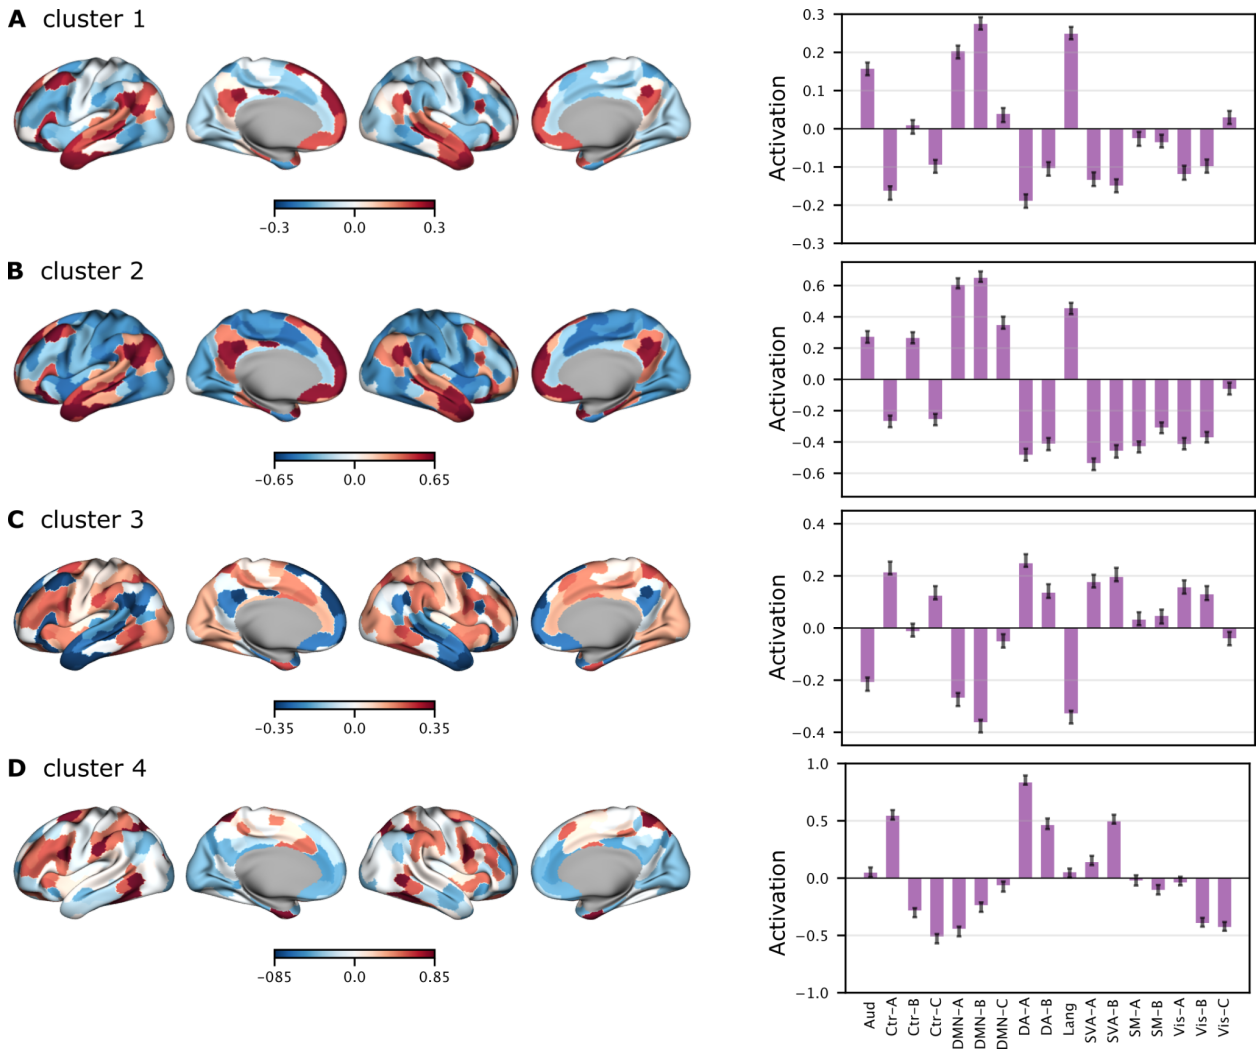

Figure 3. Brain activation patterns of representative brain states from the top four clusters. Each row shows a representative brain state from one of the top four clusters identified in the combined group HMMs. The left panel displays surface maps of average whole-brain activation, computed by averaging brain activity across all participants during timepoints when this state was active. Warm colors indicate above-average activation, and cool colors indicate below-average activation (z-scored). The right panel shows network-level activation profiles, summarizing average activation within canonical functional networks. Error bars represent the standard error across states belonging to the same cluster.

Character-specific speech modulated shared brain states

Representative brain states of clusters 1 and 2 consistently presented across narrative contexts but showed different sensitivity to specific narrative features (Figure 4). In Cluster 1,

Arthur speaking reliably increased the odds of Cluster 1 activation (OR=1.232, CI=[1.097, 1.383]), as did the presence of verbs (OR=1.153, CI=[1.032, 1.288]). Adjective usage showed reliable context differences (OR=0.87, CI=[0.769, 0.985]), indicating a stronger negative effect in the affair context (OR=0.833, P(Effect>0)=0.021) compared to paranoia (OR=1.1, P(Effect>0)=0.858). State occupancy was higher in the affair context (0.618) than in paranoia (0.538).

Cluster 2 similarly showed context interactions for Arthur-related adjectives (OR=0.797, CI=[0.594, 1.069]), with odds ratios indicating a slightly stronger negative effect in the affair context (OR=0.786, P(Effect>0)=0.128) than in paranoia (OR=0.809, P(Effect>0)=0.158). Lee and girl co-present (OR=2.42, CI=[1.817, 3.221]) and their actions (OR=1.515, CI=[1.066, 2.153]) increased brain state activation odds substantially, whereas arthur speaking (OR=0.715, CI=[0.612, 0.834]), and adverb usage (OR=0.762, CI=[0.644, 0.902]) were associated with decreases. State occupancy rate was similar in the affair context (0.121) than in paranoia (0.130).

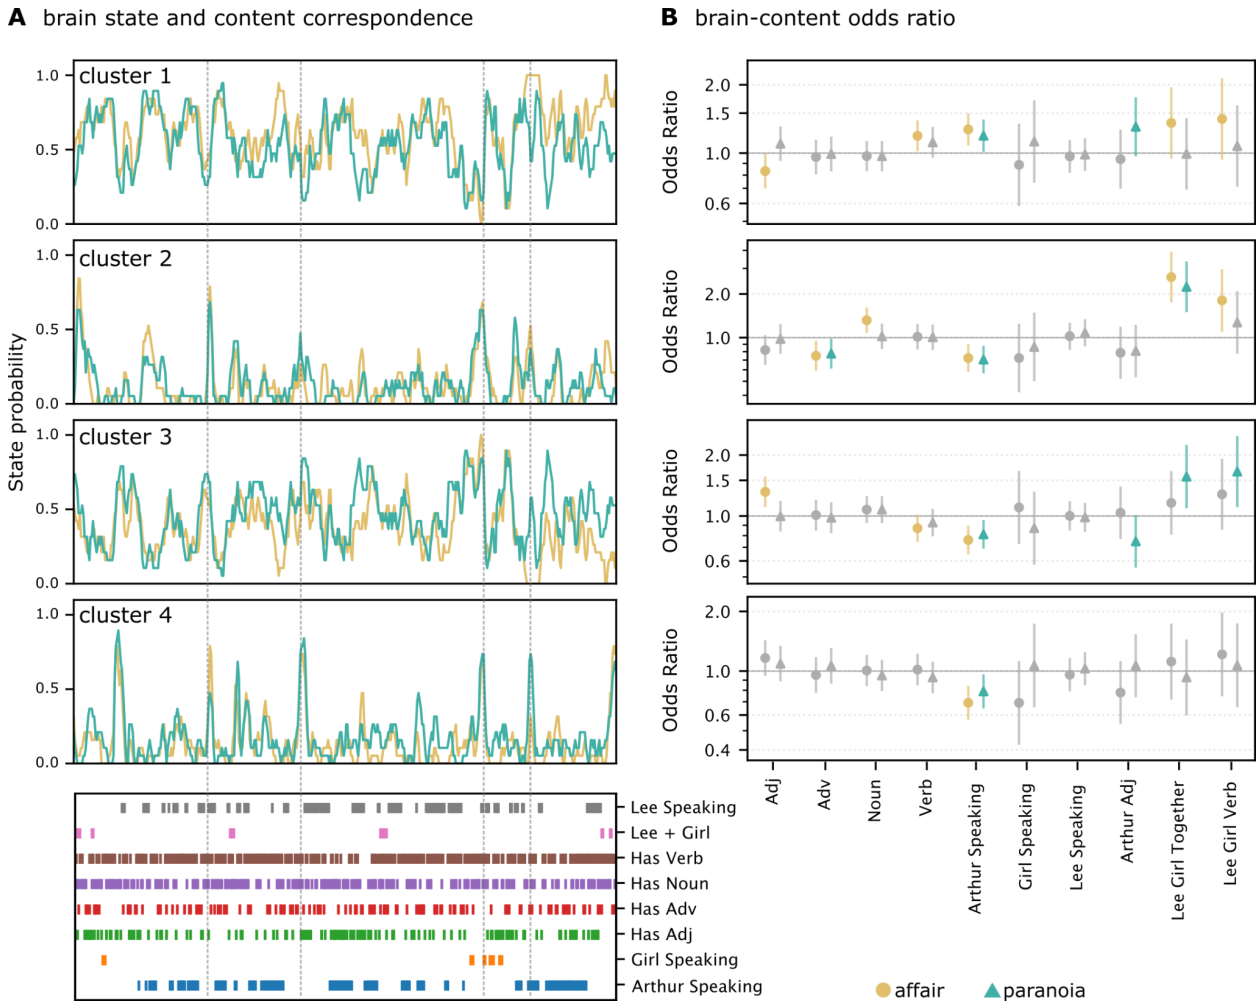

Figure 4. Correspondence between brain state dynamics and story content.
(**A**) Time series of cluster-wise brain state probabilities (gold: affair group; teal: paranoia group) for the top four HMM-derived clusters in the combined group. Each panel shows how the average probability of each brain state cluster fluctuates throughout the narrative (y-axis: state probability, range 0–1). Time series are extracted from the combined model, which includes both groups. Therefore, group differences are reflected as relative modulations in feature sensitivity (see panel B) rather than absolute occupancy. The bottom panel indicates annotated linguistic and narrative features aligned to the story timeline (x-axis represent time in TR unit), including character speech and presence, part-of-speech tags (verbs, nouns, adjectives, adverbs), and key character pairings (e.g., "Lee + Girl"). (**B**) Bayesian logistic regression analysis estimates the odds ratios (95% credible intervals) for the association between content features and brain state cluster expression. Each panel corresponds to one of the four clusters. Points represent posterior means of the odds ratio for each content predictor; error bars show 95% credible intervals. Colored markers indicate predictors with context-dependent effects that survived FDR correction in the *affair* (gold) and *paranoia* (teal) groups. Gray markers indicate predictors without statistically reliable group differences (FDR ≥ 0.05). Odds ratios above 1 suggest an increased likelihood of brain state expression when the corresponding feature is present. Conversely, odds ratios below 1 indicate a decreased likelihood of brain state expression, meaning that the feature is associated with reduced engagement of that state.

Context-specific brain states reveal diverse character influences

Representative brain states of clusters 3 and 4 exhibited distinct patterns that indicate context-specific processing mechanisms (Figure 4). Cluster 3 showed context-dependent modulation: Arthur speaking was associated with reduced odds of state activation (OR=0.786, CI=[0.7, 0.882]), while the combined presence (*lee_girl_together*, OR=1.348, CI=[1.045, 1.739]) of Lee and the girl with verbs (*lee_girl_verb*, OR=1.456, CI=[1.094, 1.938]) increased activation odds. Adjective usage increased the odds of Cluster 3 activation overall (OR=1.149, 95% CI [1.016, 1.3]), with odds ratios indicating a positive effect in the affair context (OR=1.316, P(Effect>0)=0.999) but negligible effects in paranoia (OR=0.997, P(Effect>0)=0.486). State occupancy was lower in the affair context (0.382) than paranoia (0.462).

Cluster 4 showed context interactions for Arthur speaking (OR=0.936, CI=[0.814, 1.076]). Arthur speaking demonstrated a slightly stronger negative effect in the affair context (OR=0.692, P(Effect>0)=0) compared to paranoia (OR=0.79, P(Effect>0)=0.009). Overall state occupancy was slightly lower in the affair context (0.137) than paranoia (0.171).

Multiple comparison analyses with Bayesian False Discovery Rate (FDR) corrections confirmed that the odds ratios for adjective usage (Clusters 1 and 3), Arthur-related adjectives (Clusters 1 and 2), and Arthur speaking (Cluster 4) consistently differed from chance expectations (FDR < 0.05). These results provide credible evidence that narrative context strongly modulates the temporal dynamics of brain states, revealing distinct narrative feature influences across different contexts.

Contextual modulation of behavioral responses during story comprehension

Bayesian GLMMs showed that odds ratios for participants' button presses differed systematically across contexts, indicating contextual modulation of behavioral responses.

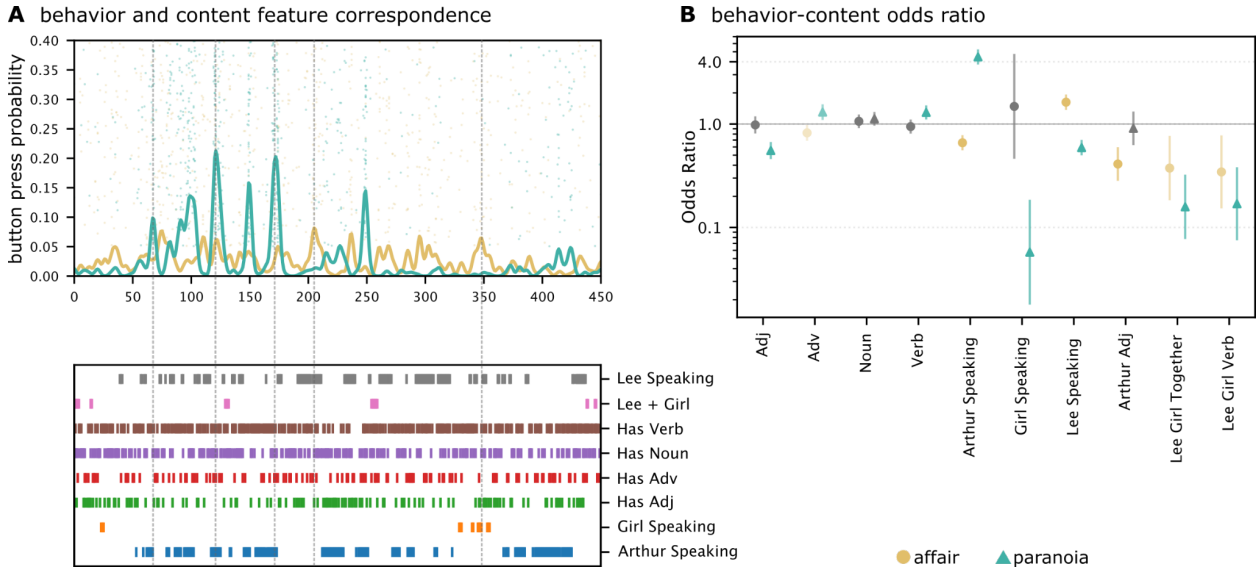

Figure 5. Correspondence between behavioral responses and story content features. (**A**) Time series of button press probabilities from participants in the *affair* (gold) and *paranoia* (teal) groups, reflecting the likelihood of detecting context-consistent narrative information over time. Dots represent individual button presses; smoothed lines indicate group-averaged response probabilities. The lower panel shows annotated story content features (identical to those in Figure 4A), including character speech, co-occurrence, and linguistic categories. (**B**) Bayesian logistic regression assessing the relationship between content features and button press behavior.

Posterior means of odds ratios (±95% credible intervals) are plotted for each predictor. Gold and teal indicate predictors with context-dependent effects that survived FDR correction in the *affair* and *paranoia* groups, respectively. Gray indicates predictors without statistically reliable group differences (FDR ≥ 0.05). Odds ratios above 1 suggest an increased likelihood of brain state expression when the corresponding feature is present. Conversely, odds ratios below 1 indicate a decreased likelihood of brain state expression, meaning that the feature is associated with reduced engagement of that state.

Context-dependent impact of character identity on behavioral responses

Character-related narrative features showed strong overall influences on behavioral responses, with Arthur speaking (OR=2.743, CI=[2.286, 3.291]) and Lee speaking (OR=1.801, CI=[1.508, 2.150]) substantially increasing the probability of button presses. Critically, context strongly modulated these effects. Arthur speaking showed large context-dependent differences in odds (interaction OR=0.334, CI=[0.279, 0.401]), with substantially greater response likelihood in the paranoia context (OR=8.205, P(Effect>0)=1.000) compared to minimal effects in the affair context (OR=0.917, P(Effect>0)=0.254). Similarly, Girl speaking was associated with higher odds in the affair context but lower odds in the paranoia context (interaction OR=2.864, CI=[1.189, 6.894]), positively influencing responses in the affair context (OR=1.559, P(Effect>0)=0.758) but strongly reducing responses in the paranoia context (OR=0.190, P(Effect>0)=0.004). Lee speaking also showed credible contextual modulation (interaction OR=0.813, CI=[0.681, 0.970]), with a stronger effect in the paranoia context (OR=2.216, P(Effect>0)>0.999) relative to the affair context (OR=1.463, P(Effect>0)=0.999).

Context-specific effects of linguistic features

Linguistic features displayed smaller credible context-dependent patterns in predicting behavioral responses. Noun usage increased the odds of button presses overall (OR=1.351, CI=[1.205, 1.515]), modulated by context (interaction OR=0.811, CI=[0.723, 0.909]). Nouns more strongly increased response probability in the paranoia context (OR=1.667, P(Effect>0)>0.999) than in the affair context (OR=1.095, P(Effect>0)=0.865). Similarly, adverb usage showed context-dependent modulation of odds (interaction OR=0.814, CI=[0.717, 0.924]), negatively influencing button presses in the affair context (OR=0.844, P(Effect>0)=0.032) but positively in paranoia (OR=1.275, P(Effect>0)=0.996). Verb usage had a modest positive main effect (OR=1.046, CI=[0.933, 1.172]) and credible context interaction (interaction OR=0.917,

CI=[0.818, 1.027]), with stronger effects observed in paranoia (OR=1.141, P(Effect>0)=0.946) than in the affair context (OR=0.959, P(Effect>0)=0.303).

Descriptive character features show contextual variation

Arthur-related adjectives were associated with reduced odds of responses overall (OR=0.411, CI=[0.297, 0.570]), with credible contextual modulation (interaction OR=1.155, CI=[0.834, 1.599]). The negative effect was stronger in the affair context (OR=0.475, P(Effect>0)=0.001) compared to paranoia (OR=0.356, P(Effect>0)<0.001). The combined presence of Lee and girl characters negatively influenced responses (OR=0.445, CI=[0.168, 1.175]), though this context modulation did not meet FDR criteria. The interaction of these characters with verbs showed no credible context-specific effects. **Supplementary Materials 2.3** confirms that the behavioral GLMM results replicate at 1-s resolution.

Discussion

Our study investigated how narrative context modulates brain state dynamics and behavioral responses during story comprehension. We identified both shared and context-specific brain states that spanned auditory, language, default mode, control, attention, and visual networks. By modeling the relationships between stimulus features, context, and brain states, we found credible evidence that context influences how narrative features, particularly speech or references to specific characters, impact the temporal dynamics of these brain states. Independent behavioral analyses revealed context-dependent differences in odds ratios for stimulus features, indicating selective modulation of behavioral responses by context.

Shared brain states suggest convergent processing during story listening

Two brain state clusters (Clusters 1 and 2) consistently appeared across both narrative contexts. Both clusters prominently activated established functional networks implicated in narrative comprehension, including auditory and language regions critical for processing spoken narratives, as well as default mode networks (DMN-A and DMN-B) associated with semantic integration, episodic memory, and narrative coherence (Hasson et al., 2018; Jackson et al., 2023; Raichle et al., 2001; Simony et al., 2016).

Interestingly, Cluster 2 included additional activation in Control Network-B and DMN-C. Control Network-B has been linked to executive function, attentional modulation, and cognitive flexibility—abilities critical for monitoring complex narrative structures and integrating detailed

episodic information (Cole et al., 2014; Dworetsky et al., 2024). DMN-C has been associated with mental simulation and contextual updating (Ritchey & Cooper, 2020). Despite shorter overall occupancy relative to Cluster 1, Cluster 2 may represent brief, intensive periods of heightened cognitive demand and/or narrative immersion.

Context-specific brain states index specialized processing demands

Clusters 3 and 4 revealed distinct brain state profiles associated with specific narrative contexts, suggesting context-dependent cognitive mechanisms. Cluster 3 was predominantly linked to the affair group and notably engaged visual processing regions, despite participants listening to an auditory stimulus. Visual network activation during story listening may represent mental imagery, facilitating visualization of narrative scenes uniquely elicited by the affair storyline (Koide-Majima et al., 2024; J. Liu et al., 2022; Pearson, 2019). Moreover, extensive activation of control, dorsal attention, and salience networks in Cluster 3 may indicate focused attentional deployment and active cognitive control processes, supporting detailed narrative evaluation and emotional engagement characteristic of the affair context (Song et al., 2023).

In contrast, Cluster 4, primarily expressed in the paranoia group, showed pronounced activity in control and salience networks. These networks have previously been associated with executive vigilance and attentional monitoring (Cole et al., 2013; Hermans et al., 2014), processes that may be particularly relevant when interpreting uncertainty, ambiguity, or potential threat—features emphasized in paranoia-inducing narratives. The co-activation of the salience network further emphasizes the importance of rapid detection and emotional appraisal of narrative cues (Lynch et al., 2024; Uddin, 2015).

The correlation structure among clusters further suggests that our results reflect a single axis of neural states spanning context-general to context-specific dynamics. This interpretation is consistent with prior work showing that low-dimensional decompositions of brain activity often yield stable states across datasets and contexts (Bolt et al., 2022; Song et al., 2023). In our case, Clusters 1 and 2 correspond to DMN-related integration states that appear across contexts, whereas Clusters 3 and 4 map onto dorsal attention and salience-related states that diverge by context. A plausible theoretical explanation is that DMN states reflect baseline integrative processes necessary for narrative comprehension, whereas attention and salience states are selectively recruited depending on whether participants adopt suspicion-related (e.g., "*Is Arthur*

paranoid?") or visualization-related (e.g., "*What are the gray-haired man and the girl doing? Are they having an affair?*") interpretive stances.

Narrative context modulates the influence of story features on brain state dynamics

Our Bayesian GLMM analyses showed that odds ratios for several narrative features differed reliably across contexts, indicating that narrative context modulates their influence on brain state dynamics. Across all clusters, brain state activations consistently responded to narrative content, underscoring a general mechanism by which narrative elements shape cognitive processing. However, notable distinctions emerged between shared clusters (Clusters 1 and 2) and context-specific clusters (Clusters 3 and 4). For instance, references to the central character, Arthur speaking, increased activation in the shared clusters but had a suppressive effect in context-specific clusters, regardless of narrative condition. The enhanced activation observed in shared clusters aligns with cognitive and psycholinguistic research demonstrating that mentions of central characters often prompt sustained attention, semantic integration, and improved memory encoding during narrative comprehension (Grall & Finn, 2022; Mar, 2011). Conversely, the suppressive effect in context-specific clusters might reflect competition between different narrative interpretations or the activation of inhibitory processes related to emotional evaluation, scenario plausibility, or context-specific narrative expectations (Sap et al., 2022; Sava-Segal et al., 2025).

Adjectives also showed contextually specific effects, aligning with prior work and indicating that the interpretation of descriptive language varies with the narrative context (Jacobs & Willems, 2018). While the precise mechanisms underlying these effects remain undetermined, the observed differences across conditions point to sensitivity in how descriptive elements are processed in context. These findings underscore the adaptability of narrative comprehension processes and their responsiveness to contextual cues embedded in story structure and content.

Contextual modulation of behavioral responses highlights selective engagement with narrative content

Independent behavioral analyses reinforce our observation that narrative context selectively modulates which stimulus features influence responses. For example, odds ratios indicated that Arthur's speaking substantially increased the likelihood of button presses in the paranoia context, but not in the affair context. This sensitivity to character speech aligns with established evidence from narrative psychology and psycholinguistics, showing heightened

632 audience engagement with central narrative figures who guide interpretive frameworks and
633 ensure story coherence (Eekhof et al., 2023; M. C. Green & Appel, 2024; Hartung et al., 2017).
634 Likewise, Lee's speaking elicited strong context-dependent behavioral responses, reflecting
635 variations in perceived narrative relevance or emotional significance of different characters
636 across contexts.

637 Importantly, these character-driven effects in behavioral responses were stronger than
638 those observed in brain-state analyses, suggesting distinct cognitive mechanisms underlying
639 explicit versus implicit narrative processing. Explicit behavioral responses likely reflect
640 deliberate inferential and evaluative processes, such as active narrative coherence assessments,
641 conscious attribution of narrative significance, or explicit inference generation. These explicit
642 processes may differ from the processes evoked by passive story listening.

643 Linguistic features such as nouns and adverbs exhibited modest but reliable context-
644 dependent effects in behavioral analyses. These differences indicate that certain lexical
645 categories may be differentially weighted or processed depending on the narrative context
646 (Tilmatine et al., 2024), though the specific mechanisms, whether semantic, evaluative, or
647 otherwise, remain to be clarified. Descriptive adjectives showed more substantial negative
648 effects in the affair context. This pattern is consistent with prior psycholinguistic findings
649 suggesting that adjectives can modulate emotional tone and guide interpretive inferences,
650 particularly in context-sensitive narratives (Lei et al., 2023).

651 These findings demonstrate that prior contextual information systematically modulated
652 behavioral responses to narrative content, with the strongest effects observed for speech
653 attribution and more modest effects for specific linguistic features. The pattern suggests that
654 context does not uniformly enhance or suppress engagement with all content features but instead
655 modulates behavioral responses in a selective, feature-dependent manner.

656 A final consideration concerns the apparent misalignment between the fMRI and
657 behavioral findings: whereas the behavioral analyses showed strong and directionally consistent
658 context effects of linguistic and character features, the fMRI-derived brain state dynamics did not
659 closely follow these behavioral directions, even though they showed statistically reliable
660 modulation. This difference reflects the fact that the two paradigms, while based on the same
661 narrative stimulus, are fundamentally different tasks. The fMRI experiment involved passive
662 listening, followed by delayed comprehension questions, which captured implicit, distributed

interpretive processes without discrete response markers. By contrast, the behavioral experiment required explicit evidence detection with overt button presses, producing stronger and more directional effects as participants actively relied on linguistic cues for decision-making. Because the tasks differ in both cognitive demands and data type (neural dynamics vs. behavioral reports), a one-to-one alignment of results is not expected. The behavioral task would be expected to engage decision- and motor-related systems not central to naturalistic comprehension. Taken together, the behavioral findings extend the fMRI results by showing how linguistic features guide explicit reasoning, while the brain-state analyses reveal implicit neural dynamics of narrative comprehension.

Limitations and future directions

Our study provides valuable insights into how narrative context modulates brain state dynamics and behavioral responses, yet several limitations point toward promising avenues for future research. First, our HMMs were conducted at the network level, averaging signals within the 17 predefined functional networks. While network-level analysis is theoretically justified given the functional specialization of these large-scale brain networks, future studies employing finer-grained analyses, such as region-of-interest or voxel-wise HMMs (Vidaurre, Abeysuriya, et al., 2018), could provide additional insights. These approaches might reveal more nuanced or heterogeneous dynamics within networks, offering a deeper understanding of localized processes during narrative comprehension.

Another limitation is that standard HMMs assume a geometric distribution of dwell times, which tends to bias results toward shorter state durations and more frequent switching between states. Hidden semi-Markov models (HSMMs) address this limitation by modeling dwell times explicitly, offering more accurate estimates of state persistence (Shappell et al., 2019). Future work could also employ hierarchical nonparametric variants such as the hierarchical Dirichlet process HMM (HDP-HMM), which infers both the number of states and their duration properties directly from the data (Beal et al., 2002; Fox et al., 2011). These approaches would provide a more flexible framework for examining naturalistic brain dynamics beyond the assumptions of conventional HMMs.

A further methodological limitation concerns model initialization. In this study, state means were initialized with random draws from a normal distribution rather than with more

structured approaches such as k-means clustering, which can sometimes improve stability for high state counts. While our validation procedures (cross-validation, bootstrap uncertainty estimation, clustering across models) minimize the risk that results depend on any single initialization scheme, future work could compare alternative initialization strategies to assess their impact on reproducibility and model fit.

Our neuroimaging and behavioral data were collected from separate participant samples. Although both datasets independently revealed context-sensitive effects, collecting brain and behavioral responses from the same individuals would allow for tighter linkage between neural state dynamics and subjective narrative judgments, enabling more direct tests of brain–behavior relationships (Xu et al., 2025).

While we used part-of-speech (PoS) labels to characterize contextual differences in language input, these labels offer only a shallow approximation of meaning. PoS categories reflect syntactic structure rather than semantic content, and future work should incorporate richer linguistic features, such as word embeddings, semantic role labels, or discourse structure, to better capture the narrative elements that drive brain dynamics.

Finally, using a single narrative stimulus may limit the generalizability of our findings to other narrative forms or genres. Future studies examining diverse narrative types, varying in emotional content, complexity, and modality, could test the breadth and boundaries of context effects on brain dynamics and behavioral engagement.

Together, these extensions would refine our understanding of how context shapes neural and behavioral responses to narrative and support more general models of naturalistic cognition.

Conclusion

This study provides converging brain and behavioral evidence that narrative context systematically modulates how listeners process and interpret story content. Through brain state modeling, we identified both context-invariant states—engaging canonical auditory, language, and default mode networks—and context-specific states that recruited control, salience, and visual networks depending on the contextual priming. These context-specific states reflected distinct patterns of functional network engagement, such as visual and attentional recruitment during the affair context, and responded differentially to narrative features, including character speech and linguistic descriptors. Complementary behavioral analyses revealed that character-

related cues, especially speech by Arthur and Lee, exerted strong and selective influences on participants' explicit judgments of narrative evidence, with the direction and magnitude of these effects shaped by contextual framing. Together, these findings demonstrate that narrative context does not merely shape interpretation retrospectively but actively reconfigures ongoing cognitive and brain processing in a temporally specific and content-sensitive manner.

Data and Code Availability

The Python code used for our analysis and visualization is available at https://github.com/yibeichan/prettymouth

Author Contributions

YC: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data Curation, Visualization, Writing - Original Draft, Project administration. ZZ: Conceptualization, Visualization, Resources, Writing - Review & Editing. SN: Conceptualization, Resources, Writing - Review & Editing, Supervision. GA: Conceptualization, Writing - Review & Editing, Supervision. SG: Funding acquisition, Resources, Writing - Review & Editing, Supervision.

Funding

YC and SSG were partially supported by NIH P4 EB019936 and by the Lann and Chris Woehrle Psychiatric Fund at the McGovern Institute for Brain Research at MIT.

Acknowledgements

We would like to thank William Menegas and the members of the Senseable Intelligence Group for insightful discussions.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Reference

Alderson-Day, B., Moffatt, J., Bernini, M., Mitrenga, K., Yao, B., & Fernyhough, C. (2020).

Processing speech and thoughts during silent reading: Direct reference effects for speech by fictional characters in voice-selective auditory cortex and a theory-of-mind network. *Journal of Cognitive Neuroscience*, *32*(9), 1637–1653. https://doi.org/10.1162/jocn_a_01571

Baker, A. P., Brookes, M. J., Rezek, I. A., Smith, S. M., Behrens, T., Probert Smith, P. J., & Woolrich, M. (2014). Fast transient networks in spontaneous human brain activity. *eLife*, *3*, e01867. https://doi.org/10.7554/eLife.01867

Baldassano, C., Chen, J., Zadbood, A., Pillow, J. W., Hasson, U., & Norman, K. A. (2017). Discovering event structure in continuous narrative perception and memory. *Neuron*, *95*(3), 709-721.e5. https://doi.org/10.1016/j.neuron.2017.06.041

Baldassano, C., Hasson, U., & Norman, K. A. (2018). Representation of real-world event schemas during narrative perception. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *38*(45), 9689–9699. https://doi.org/10.1523/JNEUROSCI.0251-18.2018

Barrett, L. F. (2022). Context reconsidered: Complex signal ensembles, relational meaning, and population thinking in psychological science. *American Psychologist*, *77*(8), 894–920. https://doi.org/10.1037/amp0001054

Beal, M. J., Ghahramani, Z., & Rasmussen, C. E. (2002). The infinite hidden Markov model. *Advances in Neural Information Processing Systems*, *14*, 577–584.

Ben-Yakov, A., Honey, C. J., Lerner, Y., & Hasson, U. (2012). Loss of reliable temporal structure in event-related averaging of naturalistic stimuli. *NeuroImage*, *63*(1), 501–506. https://doi.org/10.1016/j.neuroimage.2012.07.008

Botch, T. L., & Finn, E. S. (2024). Neural representations of concreteness and concrete concepts

are specific to the individual. *Journal of Neuroscience*, *44*(45). https://doi.org/10.1523/JNEUROSCI.0288-24.2024

Bransford, J. D., & Johnson, M. K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior*, *11*(6), 717–726. https://doi.org/10.1016/S0022-5371(72)80006-9

Coderre, E. L., & Cohn, N. (2023). Individual differences in the neural dynamics of visual narrative comprehension: The effects of proficiency and age of acquisition. *Psychonomic Bulletin & Review*. https://doi.org/10.3758/s13423-023-02334-x

Cole, M. W., Bassett, D. S., Power, J. D., Braver, T. S., & Petersen, S. E. (2014). Intrinsic and task-evoked network architectures of the human brain. *Neuron*, *83*(1), 238–251. https://doi.org/10.1016/j.neuron.2014.05.014

Cole, M. W., Reynolds, J. R., Power, J. D., Repovs, G., Anticevic, A., & Braver, T. S. (2013). Multi-task connectivity reveals flexible hubs for adaptive task control. *Nature Neuroscience*, *16*(9), 1348–1355. https://doi.org/10.1038/nn.3470

de Bruin, D., van Baar, J. M., Rodríguez, P. L., & FeldmanHall, O. (2023). Shared neural representations and temporal segmentation of political content predict ideological similarity. *Science Advances*, *9*(5), eabq5920. https://doi.org/10.1126/sciadv.abq5920

Dworetsky, A., Seitzman, B. A., Adeyemo, B., Nielsen, A. N., Hatoum, A. S., Smith, D. M., Nichols, T. E., Neta, M., Petersen, S. E., & Gratton, C. (2024). Two common and distinct forms of variation in human functional brain networks. *Nature Neuroscience*, *27*(6), 1187–1198. https://doi.org/10.1038/s41593-024-01618-2

Eekhof, L. S., Krieken, K. van, Sanders, J., & Willems, R. M. (2023). Engagement with narrative characters: The role of social-cognitive abilities and linguistic viewpoint. *Discourse*

Processes. https://www.tandfonline.com/doi/abs/10.1080/0163853X.2023.2206773

Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., Kent, J. D., Goncalves, M., DuPre, E., Snyder, M., Oya, H., Ghosh, S. S., Wright, J., Durnez, J., Poldrack, R. A., & Gorgolewski, K. J. (2019). fMRIPrep: A robust preprocessing pipeline for functional MRI. *Nature Methods*, *16*(1), Article 1. https://doi.org/10.1038/s41592-018-0235-4

Fox, E. B., Sudderth, E. B., Jordan, M. I., & Willsky, A. S. (2011). A sticky HDP-HMM with application to speaker diarization. *Annals of Applied Statistics, 5*(2A), 1020–1056. https://doi.org/10.1214/10-AOAS395

Geerligs, L., Gözükara, D., Oetringer, D., Campbell, K. L., van Gerven, M., & Güçlü, U. (2022). A partially nested cortical hierarchy of neural states underlies event segmentation in the human brain. *eLife*, *11*, e77430. https://doi.org/10.7554/eLife.77430

Gerrig, R. J. (1993). *Experiencing narrative worlds: On the psychological activities of reading*. Routledge. https://doi.org/10.4324/9780429500633

Grall, C., & Finn, E. S. (2022). Leveraging the power of media to drive cognition: A media-informed approach to naturalistic neuroscience. *Social Cognitive and Affective Neuroscience*, nsac019. https://doi.org/10.1093/scan/nsac019

Green, M. C., & Appel, M. (2024). Chapter One - Narrative transportation: How stories shape how we see ourselves and the world. In B. Gawronski (Ed.), *Advances in Experimental Social Psychology* (Vol. 70, pp. 1–82). Academic Press. https://doi.org/10.1016/bs.aesp.2024.03.002

Green, M. J., & Phillips, M. L. (2004). Social threat perception and the evolution of paranoia. *Neuroscience & Biobehavioral Reviews*, *28*(3), 333–342.

https://doi.org/10.1016/j.neubiorev.2004.03.006

Hartung, F., Withers, P., Hagoort, P., & Willems, R. M. (2017). When fiction is just as real as fact: No differences in reading behavior between stories believed to be based on true or fictional events. *Frontiers in Psychology*, *8*. https://doi.org/10.3389/fpsyg.2017.01618

Hasson, U., Chen, J., & Honey, C. J. (2015). Hierarchical process memory: Memory as an integral component of information processing. *Trends in Cognitive Sciences*, *19*(6), 304–313. https://doi.org/10.1016/j.tics.2015.04.006

Hasson, U., Egidi, G., Marelli, M., & Willems, R. M. (2018). Grounding the neurobiology of language in first principles: The necessity of non-language-centric explanations for language comprehension. *Cognition*, *180*, 135–157. https://doi.org/10.1016/j.cognition.2018.06.018

Hermans, E. J., Henckens, M. J. A. G., Joëls, M., & Fernández, G. (2014). Dynamic adaptation of large-scale brain networks in response to acute stressors. *Trends in Neurosciences*, *37*(6), 304–314. https://doi.org/10.1016/j.tins.2014.03.006

Jackson, R. L., Humphreys, G. F., Rice, G. E., Binney, R. J., & Lambon Ralph, M. A. (2023). A network-level test of the role of the co-activated default mode network in episodic recall and social cognition. *Cortex: A Journal Devoted to the Study of the Nervous System and Behavior*, *165*, 141–159. https://doi.org/10.1016/j.cortex.2022.12.016

Jacobs, A. M., & Willems, R. M. (2018). The fictive brain: Neurocognitive correlates of engagement in literature. *Review of General Psychology*, *22*(2), 147–160. https://doi.org/10.1037/gpr0000106

Jacoby, N., & Fedorenko, E. (2020). Discourse-level comprehension engages medial frontal Theory of Mind brain regions even for expository texts. *Language, Cognition and*

Neuroscience, *35*(6), 780–796. https://doi.org/10.1080/23273798.2018.1525494

Johns, C. L., Jahn, A. A., Jones, H. R., Kush, D., Molfese, P. J., Van Dyke, J. A., Magnuson, J. S., Tabor, W., Mencl, W. E., Shankweiler, D. P., & Braze, D. (2018). Individual differences in decoding skill, print exposure, and cortical structure in young adults. *Language, Cognition and Neuroscience*, *33*(10), 1275–1295. https://doi.org/10.1080/23273798.2018.1476727

Koide-Majima, N., Nishimoto, S., & Majima, K. (2024). Mental image reconstruction from human brain activity: Neural decoding of mental imagery via deep neural network-based Bayesian estimation. *Neural Networks*, *170*, 349–363. https://doi.org/10.1016/j.neunet.2023.11.024

Kong, R., Yang, Q., Gordon, E., Xue, A., Yan, X., Orban, C., Zuo, X.-N., Spreng, N., Ge, T., Holmes, A., Eickhoff, S., & Yeo, B. T. T. (2021). Individual-specific areal-level parcellations improve functional connectivity prediction of behavior. *Cerebral Cortex*, *31*(10), 4477–4500. https://doi.org/10.1093/cercor/bhab101

Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, *2*(1–2), 83–97. https://doi.org/10.1002/nav.3800020109

Lahnakoski, J. M., Glerean, E., Jääskeläinen, I. P., Hyönä, J., Hari, R., Sams, M., & Nummenmaa, L. (2014). Synchronous brain activity across individuals underlies shared psychological perspectives. *NeuroImage*, *100*, 316–324. https://doi.org/10.1016/j.neuroimage.2014.06.022

Lei, A., Willems ,Roel M., & and Eekhof, L. S. (2023). Emotions, fast and slow: Processing of emotion words is affected by individual differences in need for affect and narrative absorption. *Cognition and Emotion*, *37*(5), 997–1005.

https://doi.org/10.1080/02699931.2023.2216445

Lerner, Y., Honey, C. J., Silbert, L. J., & Hasson, U. (2011). Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *Journal of Neuroscience*, *31*(8), 2906–2915. https://doi.org/10.1523/JNEUROSCI.3684-10.2011

Li, A., Liu, H., Lei, X., He, Y., Wu, Q., Yan, Y., ... & Liu, B. (2023). Hierarchical fluctuation shapes a dynamic flow linked to states of consciousness. *Nature communications, 14*(1), 3238. https://doi.org/10.1038/s41467-023-38972-x

Liu, J., Spagna, A., & Bartolomeo, P. (2022). Hemispheric asymmetries in visual mental imagery. *Brain Structure & Function*, *227*(2), 697–708. https://doi.org/10.1007/s00429-021-02277-w

Liu, L., Jiang, J., Li, H., & Ding, G. (2025). Tripartite organization of brain state dynamics underlying spoken narrative comprehension. *eLife*, *13*. https://doi.org/10.7554/eLife.99997.2

Lynch, C. J., Elbau, I. G., Ng, T., Ayaz, A., Zhu, S., Wolk, D., Manfredi, N., Johnson, M., Chang, M., Chou, J., Summerville, I., Ho, C., Lueckel, M., Bukhari, H., Buchanan, D., Victoria, L. W., Solomonov, N., Goldwaser, E., Moia, S., … Liston, C. (2024). Frontostriatal salience network expansion in individuals in depression. *Nature*, *633*(8030), 624–633. https://doi.org/10.1038/s41586-024-07805-2

Mar, R. A. (2011). The neural bases of social cognition and story comprehension. *Annual Review of Psychology*, *62*(Volume 62, 2011), 103–134. https://doi.org/10.1146/annurev-psych-120709-145406

Mar, R. A., & Oatley, K. (2008). The function of fiction is the abstraction and simulation of social experience. *Perspectives on Psychological Science*, *3*(3), 173–192.

https://doi.org/10.1111/j.1745-6924.2008.00073.x

McIntosh, A. R. (2004). Contexts and catalysts. *Neuroinformatics*, *2*(2), 175–181. https://doi.org/10.1385/NI:2:2:175

Meer, J. N. van der, Breakspear, M., Chang, L. J., Sonkusare, S., & Cocchi, L. (2020). Movie viewing elicits rich and reliable brain state dynamics. *Nature Communications*, *11*(1), Article 1. https://doi.org/10.1038/s41467-020-18717-w

Nastase, S. A., Liu, Y.-F., Hillman, H., Zadbood, A., Hasenfratz, L., Keshavarzian, N., Chen, J., Honey, C. J., Yeshurun, Y., Regev, M., Nguyen, M., Chang, C. H. C., Baldassano, C., Lositsky, O., Simony, E., Chow, M. A., Leong, Y. C., Brooks, P. P., Micciche, E., … Hasson, U. (2021). The "Narratives" fMRI dataset for evaluating models of naturalistic language comprehension. *Scientific Data*, *8*(1), 250. https://doi.org/10/gnpts8

Nguyen, M., Vanderwal, T., & Hasson, U. (2019). Shared understanding of narratives is correlated with shared neural responses. *NeuroImage*, *184*, 161–170. https://doi.org/10.1016/j.neuroimage.2018.09.010

Nieuwland, M. S., & Van Berkum, J. J. A. (2006). When peanuts fall in love: N400 evidence for the power of discourse. *Journal of Cognitive Neuroscience*, *18*(7), 1098–1111. https://doi.org/10.1162/jocn.2006.18.7.1098

Nijhof, A. D., & Willems, R. M. (2015). Simulating fiction: Individual differences in literature comprehension revealed with fMRI. *PLOS ONE*, *10*(2), e0116492. https://doi.org/10.1371/journal.pone.0116492

Pearson, J. (2019). The human imagination: The cognitive neuroscience of visual mental imagery. *Nature Reviews Neuroscience*, *20*(10), 624–634. https://doi.org/10.1038/s41583-019-0202-9

Quinn, A. J., Vidaurre, D., Abeysuriya, R., Becker, R., Nobre, A. C., & Woolrich, M. W. (2018). Task-evoked dynamic network analysis through hidden markov modeling. *Frontiers in Neuroscience*, *12*, 603. https://doi.org/10.3389/fnins.2018.00603

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences*, *98*(2), 676–682. https://doi.org/10.1073/pnas.98.2.676

Ritchey, M., & Cooper, R. A. (2020). Deconstructing the posterior medial episodic network. *Trends in Cognitive Sciences*, *24*(6), 451–465. https://doi.org/10.1016/j.tics.2020.03.006

Sap, M., Jafarpour, A., Choi, Y., Smith, N. A., Pennebaker, J. W., & Horvitz, E. (2022). Quantifying the narrative flow of imagined versus autobiographical stories. *Proceedings of the National Academy of Sciences*, *119*(45), e2211715119. https://doi.org/10.1073/pnas.2211715119

Sava-Segal, C., Grall, C., & Finn, E. S. (2025). *Narrative 'twist' shifts within-individual neural representations of dissociable story features* (p. 2025.01.13.632631). bioRxiv. https://doi.org/10.1101/2025.01.13.632631

Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., Eickhoff, S. B., & Yeo, B. T. T. (2018). Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral Cortex*, *28*(9), 3095–3114. https://doi.org/10.1093/cercor/bhx179

Shappell, H. M., Caffo, B. S., Pekar, J. J., & Lindquist, M. A. (2019). Improved state change estimation in dynamic functional connectivity using hidden semi-Markov models. *NeuroImage*, *191*, 243–257. https://doi.org/10.1016/j.neuroimage.2019.02.013

Shine, J. M., Breakspear, M., Bell, P. T., Ehgoetz Martens, K. A., Shine, R., Koyejo, O., Sporns,

O., & Poldrack, R. A. (2019). Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. *Nature Neuroscience*, *22*(2), 289–296. https://doi.org/10.1038/s41593-018-0312-0

Simony, E., Honey, C. J., Chen, J., Lositsky, O., Yeshurun, Y., Wiesel, A., & Hasson, U. (2016). Dynamic reconfiguration of the default mode network during narrative comprehension. *Nature Communications*, *7*(1), 12141. https://doi.org/10.1038/ncomms12141

Song, H., Park, B., Park, H., & Shim, W. M. (2021). Cognitive and neural state dynamics of narrative comprehension. *Journal of Neuroscience*, *41*(43), 8972–8990. https://doi.org/10.1523/JNEUROSCI.0037-21.2021

Song, H., Shim, W. M., & Rosenberg, M. D. (2023). Large-scale neural dynamics in a shared low-dimensional state space reflect cognitive and attentional dynamics. *eLife*, *12*, e85487. https://doi.org/10.7554/eLife.85487

Taghia, J., Cai, W., Ryali, S., Kochalka, J., Nicholas, J., Chen, T., & Menon, V. (2018). Uncovering hidden brain state dynamics that regulate performance and decision-making during cognition. *Nature Communications*, *9*(1), 2505. https://doi.org/10.1038/s41467-018-04723-6

Tilmatine, M., Lüdtke, J., & Jacobs, A. M. (2024). Predicting subjective ratings of affect and comprehensibility with text features: A reader response study of narrative poetry. *Frontiers in Psychology*, *15*. https://doi.org/10.3389/fpsyg.2024.1431764

Uddin, L. Q. (2015). Salience processing and insular cortical function and dysfunction. *Nature Reviews Neuroscience*, *16*(1), 55–61. https://doi.org/10.1038/nrn3857

van Kesteren, M. T. R., Beul, S. F., Takashima, A., Henson, R. N., Ruiter, D. J., & Fernández, G. (2013). Differential roles for medial prefrontal and medial temporal cortices in schema-

dependent encoding: From congruent to incongruent. *Neuropsychologia*, *51*(12), 2352–2359. https://doi.org/10.1016/j.neuropsychologia.2013.05.027

Vidaurre, D., Abeysuriya, R., Becker, R., Quinn, A. J., Alfaro-Almagro, F., Smith, S. M., & Woolrich, M. W. (2018). Discovering dynamic brain networks from big data in rest and task. *Neuroimage*, *180*(Pt B), 646–656. https://doi.org/10.1016/j.neuroimage.2017.06.077

Vidaurre, D., Hunt, L. T., Quinn, A. J., Hunt, B. A. E., Brookes, M. J., Nobre, A. C., & Woolrich, M. W. (2018). Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks. *Nature Communications*, *9*(1), 2987. https://doi.org/10.1038/s41467-018-05316-z

Vidaurre, D., Smith, S. M., & Woolrich, M. W. (2017). Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*, *114*(48), 12827–12832. https://doi.org/10.1073/pnas.1705120114

Willems, R. M., Nastase, S. A., & Milivojevic, B. (2020). Narratives for neuroscience. *Trends in Neurosciences*, *43*(5), 271–273. https://doi.org/10.1016/j.tins.2020.03.003

Xu, T., Kiar, G., Zuo, X.-N., Vogelstein, J. T., & Milham, M. P. (2025). Challenges in measuring individual differences of brain function. *Imaging Neuroscience*, *3*, imag_a_00430. https://doi.org/10.1162/imag_a_00430

Yang, E., Milisav, F., Kopal, J., Holmes, A. J., Mitsis, G. D., Misic, B., Finn, E. S., & Bzdok, D. (2023). The default network dominates neural responses to evolving movie stories. *Nature Communications*, *14*(1), 4197. https://doi.org/10.1038/s41467-023-39862-y

Yarkoni, T., Speer, N. K., Balota, D. A., McAvoy, M. P., & Zacks, J. M. (2008). Pictures of a thousand words: Investigating the neural mechanisms of reading with extremely rapid event-related fMRI. *NeuroImage*, *42*(2), 973–987.

https://doi.org/10.1016/j.neuroimage.2008.04.258

Yeshurun, Y., Nguyen, M., & Hasson, U. (2021). The default mode network: Where the idiosyncratic self meets the shared social world. *Nature Reviews Neuroscience*, *22*(3), 181–192. https://doi.org/10.1038/s41583-020-00420-w

Yeshurun, Y., Swanson, S., Simony, E., Chen, J., Lazaridi, C., Honey, C. J., & Hasson, U. (2017). Same story, different story: The neural representation of interpretive frameworks. *Psychological Science*, *28*(3), 307–319. https://doi.org/10.1177/0956797616682029

Zhao, C., Jarecka, D., Covitz, S., Chen, Y., Eickhoff, S. B., Fair, D. A., Franco, A. R., Halchenko, Y. O., Hendrickson, T. J., Hoffstaedter, F., Houghton, A., Kiar, G., Macdonald, A., Mehta, K., Milham, M. P., Salo, T., Hanke, M., Ghosh, S. S., Cieslak, M., & Satterthwaite, T. D. (2024). A reproducible and generalizable software workflow for analysis of large-scale neuroimaging data collections using BIDS Apps. *Imaging Neuroscience*, *2*, 1–19. https://doi.org/10.1162/imag_a_00074

Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, *123*(2), 162–185. https://doi.org/10.1037/0033-2909.123.2.162

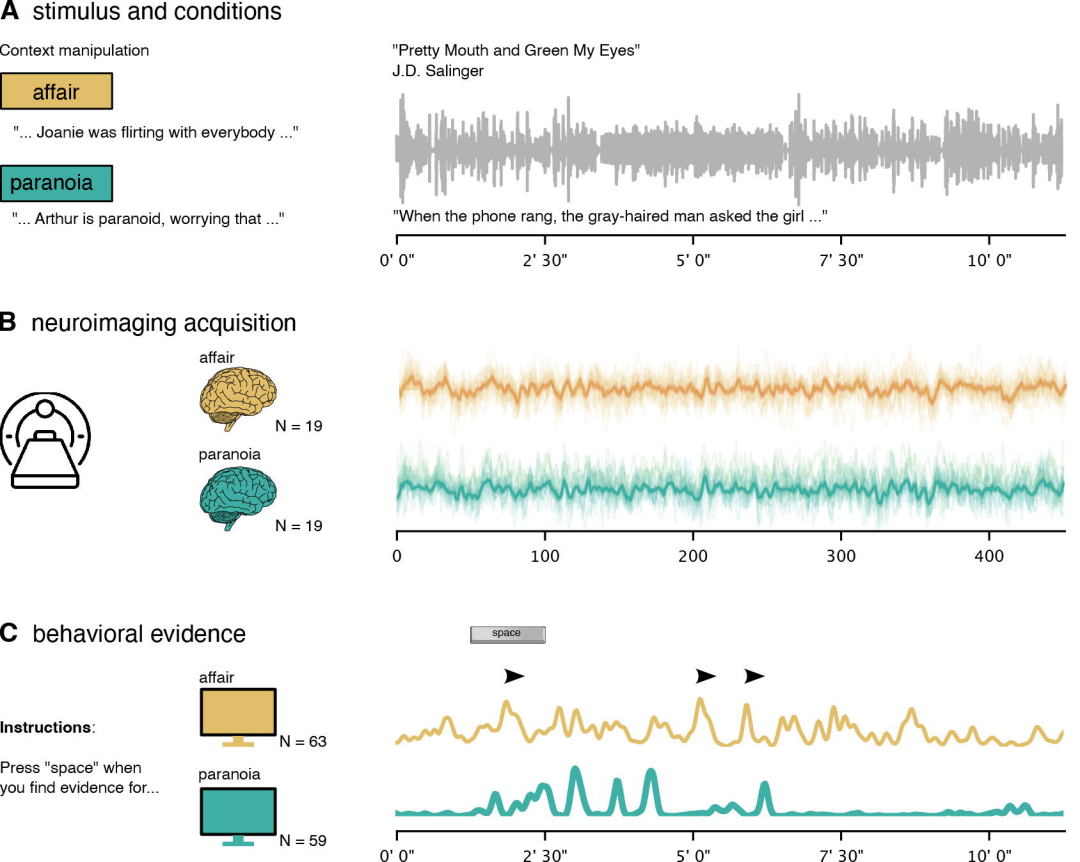
A stimulus and conditions
Context manipulation
affair
"... Joanie was flirting with everybody ..."
paranoia
"... Arthur is paranoid, worrying that ..."
"Pretty Mouth and Green My Eyes"
J.D. Salinger
"When the phone rang, the gray-haired man asked the girl ..."
0' 0"
2' 30"
5' 0"
7' 30"
10' 0"
B neuroimaging acquisition
affair
N = 19
paranoia
N = 19
0
100
200
300
400
C behavioral evidence
Instructions:
Press "space" when you find evidence for...
affair
N = 63
paranoia
N = 59
space
0' 0"
2' 30"
5' 0"
7' 30"
10' 0"

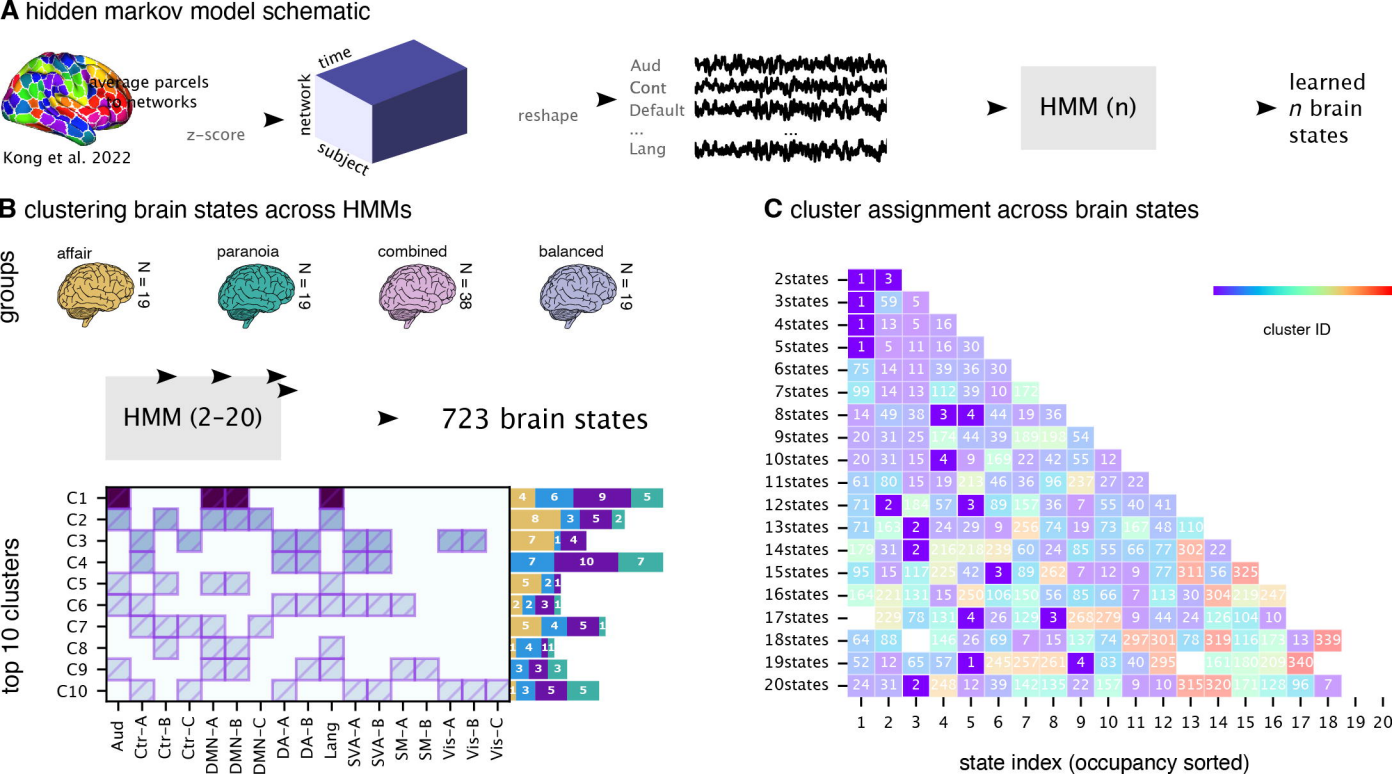
A hidden markov model schematic
average parcels to networks
Kong et al. 2022
z-score
network
time
subject
reshape
Aud
Cont
Default
...
Lang
HMM (n)
learned n brain states
B clustering brain states across HMMs
groups
affair
N = 19
paranoia
N = 19
combined
N = 38
balanced
N = 19
HMM (2-20)
723 brain states
top 10 clusters
C1
C2
C3
C4
C5
C6
C7
C8
C9
C10
Aud
Ctr-A
Ctr-B
Ctr-C
DMN-A
DMN-B
DMN-C
DA-A
DA-B
Lang
SVA-A
SVA-B
SM-A
SM-B
Vis-A
Vis-B
Vis-C
C cluster assignment across brain states
2states
3states
4states
5states
6states
7states
8states
9states
10states
11states
12states
13states
14states
15states
16states
17states
18states
19states
20states
cluster ID
state index (occupancy sorted)

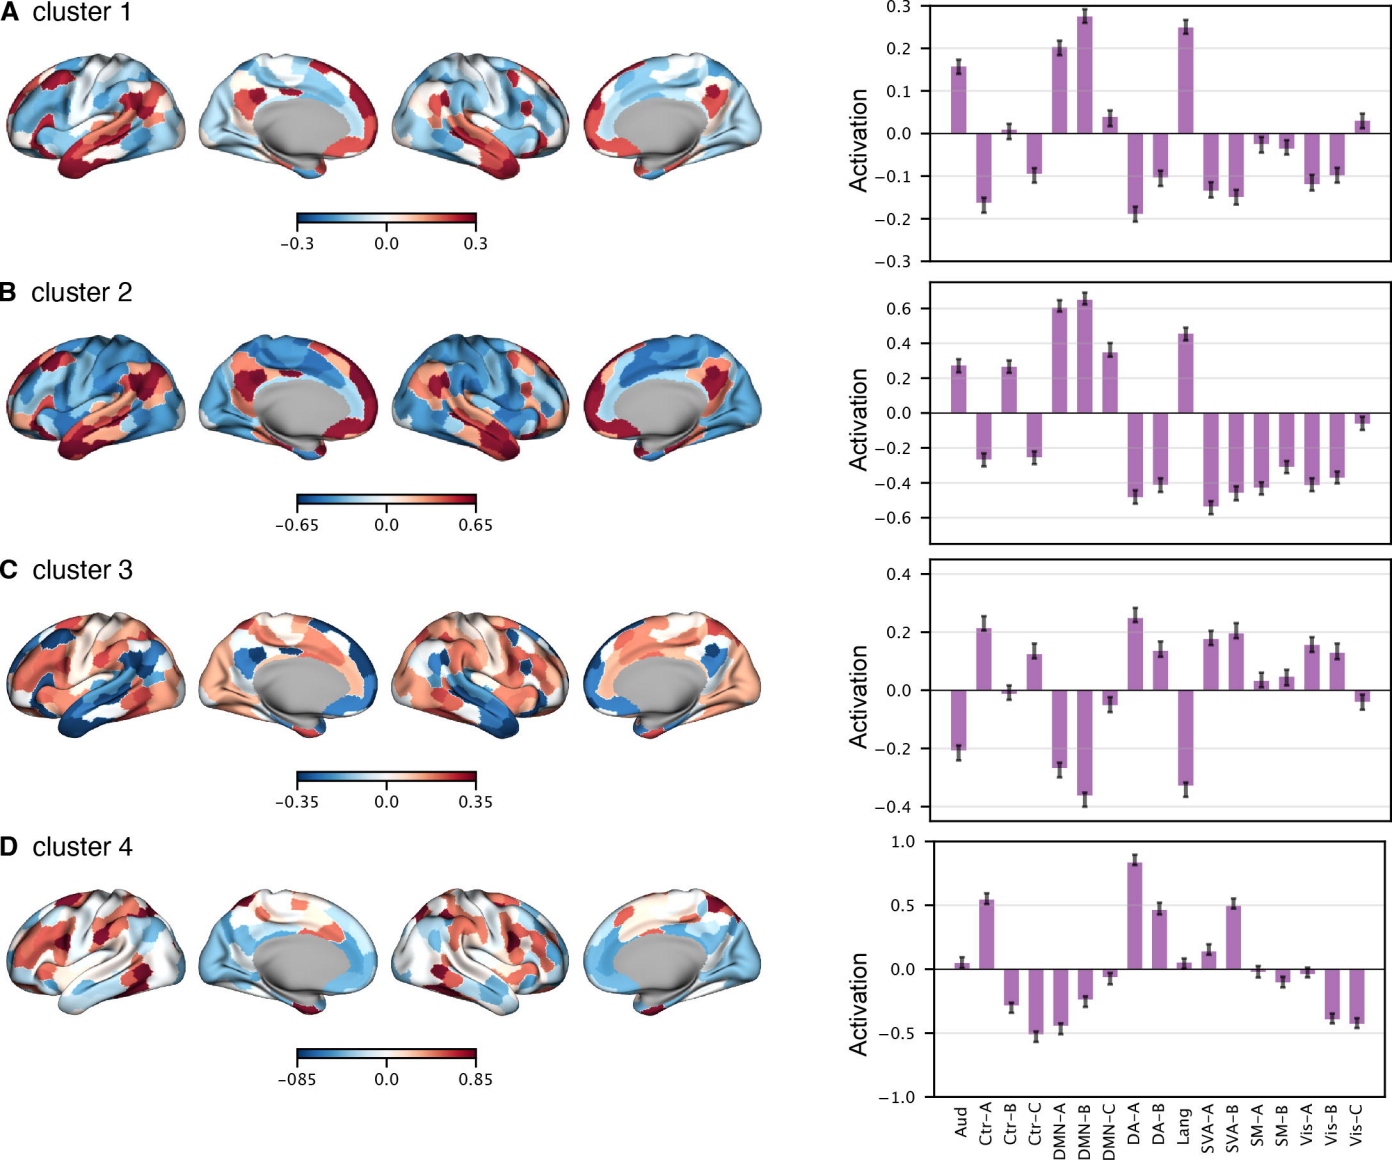
A cluster 1
-0.3
0.0
0.3
Activation
B cluster 2
-0.65
0.0
0.65
Activation
C cluster 3
-0.35
0.0
0.35
Activation
D cluster 4
-085
0.0
0.85
Activation
Aud
Ctr-A
Ctr-B
Ctr-C
DMN-A
DMN-B
DMN-C
DA-A
DA-B
Lang
SVA-A
SVA-B
SM-A
SM-B
Vis-A
Vis-B
Vis-C

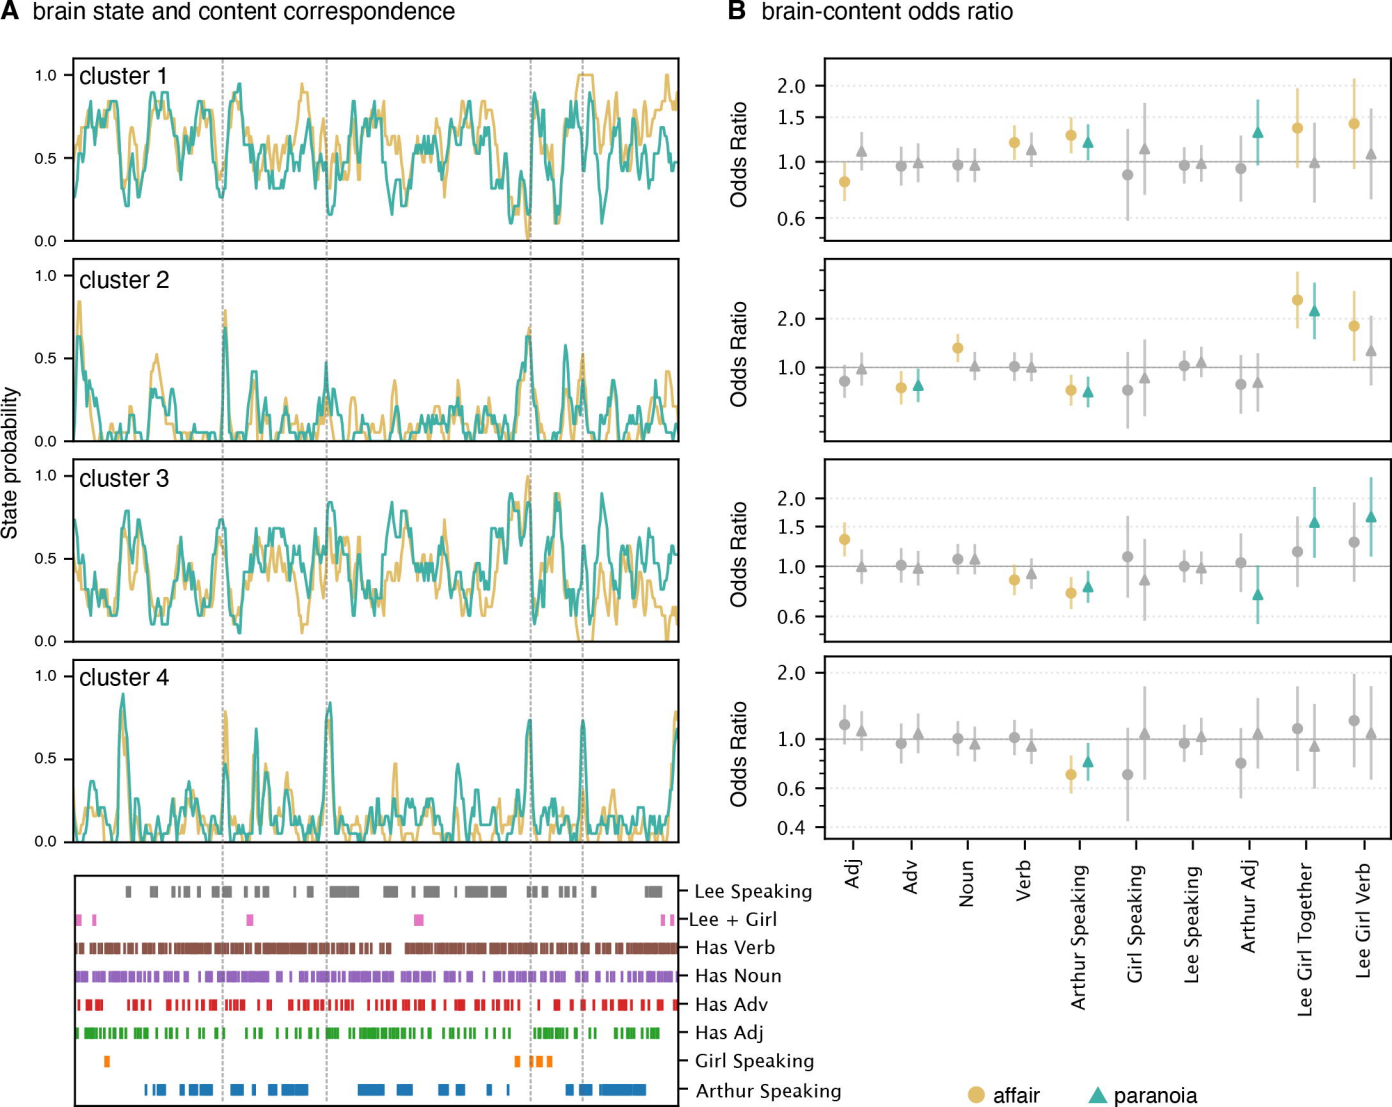
A brain state and content correspondence
B brain-content odds ratio
cluster 1
cluster 2
cluster 3
cluster 4
State probability
1.0
0.5
0.0
Lee Speaking
Lee + Girl
Has Verb
Has Noun
Has Adv
Has Adj
Girl Speaking
Arthur Speaking
Odds Ratio
2.0
1.5
1.0
0.6
0.4
Adj
Adv
Noun
Verb
Arthur Speaking
Girl Speaking
Lee Speaking
Arthur Adj
Lee Girl Together
Lee Girl Verb
affair
paranoia

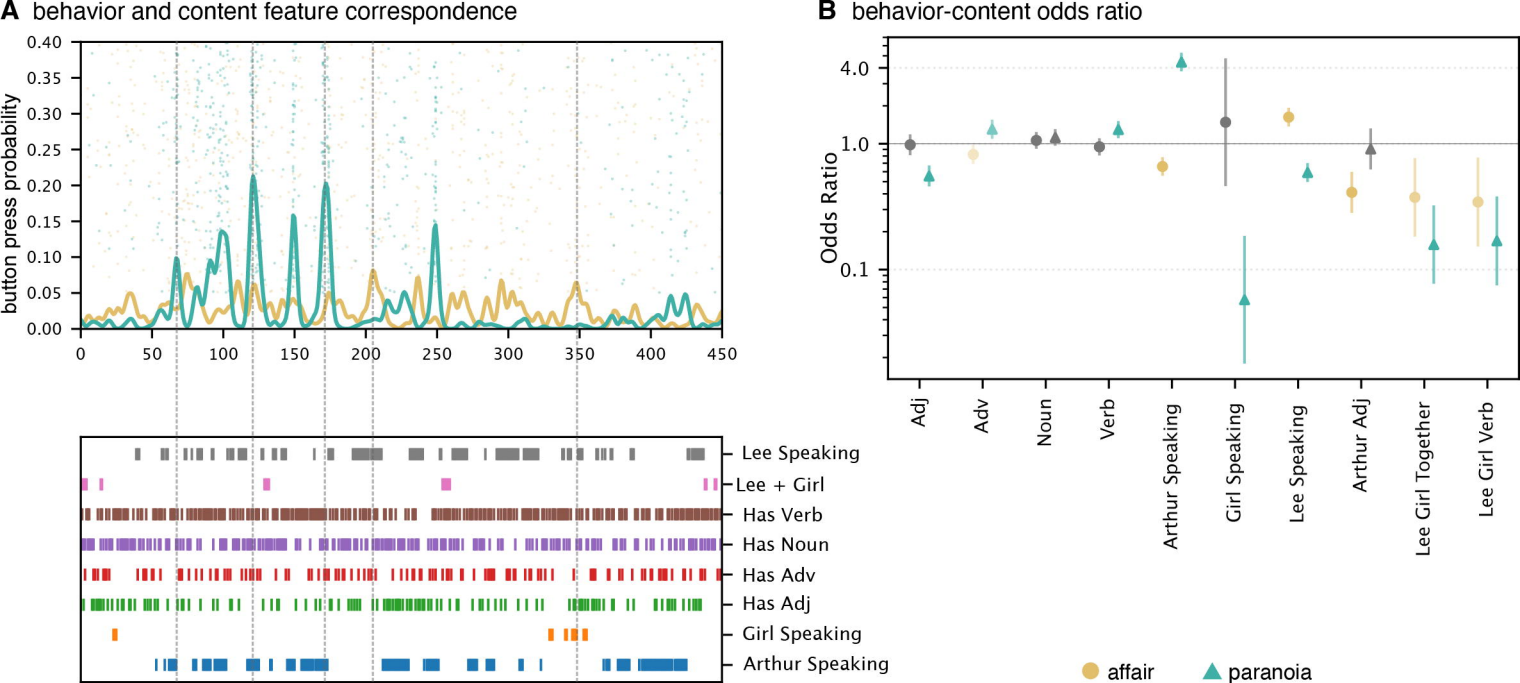
A behavior and content feature correspondence
button press probability
0.40
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
0
50
100
150
200
250
300
350
400
450
Lee Speaking
Lee + Girl
Has Verb
Has Noun
Has Adv
Has Adj
Girl Speaking
Arthur Speaking
B behavior-content odds ratio
Odds Ratio
4.0
1.0
0.1
Adj
Adv
Noun
Verb
Arthur Speaking
Girl Speaking
Lee Speaking
Arthur Adj
Lee Girl Together
Lee Girl Verb
affair
paranoia