Anticipation of temporally structured events in the brain

Caroline S. Lee[1,2], Mariam Aly[1], and Christopher Baldassano[1]

[1]Columbia University, Department of Psychology
[2]Dartmouth College, Department of Psychological and Brain Sciences

1 Abstract

Higher-order regions in the brain play a role in organizing perceptual input into semantically-meaningful schematic representations, and they do so at differing timescales up the cortical hierarchy. Such integration in higher-order regions critically depends on the availability of temporal structure over long timescales. The ability to extract temporal structure is particularly adaptive because it enables the generation of predictions about upcoming events. How does the brain use repeated experience in structured environments to anticipate what is likely to come next? Here, we examine how the brain can anticipate event boundaries in familiar sequences of actions. We used a naturalistic stimulus (a movie), in which regularities are present at multiple timescales, enabling us to identify multiple timescales of prediction in the brain. We hypothesized that the timescale of neural prediction would vary continuously, with progressively higher-order regions (e.g., prefrontal cortex) predicting further in the future than lower-order regions (e.g., visual cortex). To test this, we examined brain activity with fMRI while individuals watched a 90-second clip from the movie The Grand Budapest Hotel, six times. To uncover neural anticipation, we used a searchlight approach in which, for each region of interest, we fit a Hidden Markov Model to identify temporal shifts between multivariate activity patterns evoked by the first viewing of the movie clip compared to repeated viewings. We found temporal shifts in event patterns throughout the brain, such that neural activity patterns for repeated viewings shifted to preceded those of initial viewing. The scale of prediction varied hierarchically in a posterior (less prediction) to anterior (more prediction) fashion, with prediction reaching up to 12 seconds in the future. Furthermore, in a subset of regions exhibiting prediction, neural event boundaries shifted with repeated viewing to precede subjective event boundaries by 6-7 seconds. Together, these results demonstrate that prediction is ubiquitous in the brain, but the amount of prediction varies across the cortical hierarchy.

2 Introduction

A primary function of the brain is to adaptively use past experience to generate predictions about events that are likely to occur in the future (Clark, 2013; Friston, 2005). Such prediction is ubiquitous in the brain, spanning systems that support sensation, action, memory, motivation, and language (den Ouden, Daunizeau, Roiser, Friston, & Stephan, 2010). For example, the visual system takes advantage of the world's relative stability over space and time to anticipate upcoming input (de Lange, Heilbron, & Kok, 2018; Summerfield & Egner, 2009). The majority of studies of prediction, however, have tested only one-step associations between pairs of simple and static stimuli, such as lines, dots, oriented gratings, or abstract objects (e.g., Alink, Schwiedrzik, Kohler, Singer, & Muckli, 2010; Ekman, Kok, & de Lange, 2017; Gavornik & Bear, 2014; Hindy, Ng, & Turk-Browne, 2016; Kok, Jehee, & de Lange, 2012; Kok, Failing, & de Lange, 2014). These studies have found predictive signals about a single upcoming stimulus in a variety of brain regions, from perceptual regions (Kok et al., 2012, 2014) to the medial temporal lobe (Hindy et al., 2016). How does the brain use repeated experience in naturalistic environments to anticipate upcoming sequences of events that extend farther into the future?

Prior work has shown that the brain integrates information about the recent past over a hierarchy of timescales (Aly, Chen, Turk-Browne, & Hasson, 2018; Hasson, Chen, & Honey, 2015). Lower-order areas primarily represent the current moment, whereas higher-order areas are sensitive to information from many seconds or even minutes into the past. Higher-order regions with longer timescales play a critical role in organizing perceptual input into semantically-meaningful schematic representations (Baldassano et al., 2017; Baldassano, Hasson, & Norman, 2018). What is less clear is whether this hierarchy also exists for prediction: as we move from lower-order perceptual systems into higher-order areas, do these regions exhibit different timescales of predictions about the future? We previously found that higher-order regions did exhibit anticipatory signals when individuals had prior knowledge of the general structure of a narrative (Baldassano et al., 2018). But these individuals only had knowledge of information at relatively long timescales (e.g., the general sequence of events, and not moment-by-moment perceptual features), so we were unable to assess whether they could generate predictions across the timescale hierarchy.

Here, we examine how the brain anticipates event boundaries in familiar sequences of actions. We used a naturalistic narrative stimulus (a movie), in which regularities are present at multiple timescales. For example, upon second viewing of a movie, one can anticipate the next action to be taken in a given scene, the next character to appear, the next location that is visited, and the last scene of the movie. The presence of predictability at multiple timescales in the same stimulus enables us to identify varying timescales of prediction in the brain that co-exist simultaneously. We hypothesized that the timescale of neural prediction would vary continuously, with progressively higher-order regions (e.g., prefrontal cortex) predicting further in the future than lower-order regions (e.g., visual cortex).

To test this, we examined brain activity with fMRI while individuals watched a 90-second clip from the movie *The Grand Budapest Hotel*, 6 times. To uncover neural anticipation, we used a searchlight approach in which, for each region of interest, we fit a Hidden Markov Model (HMM) to identify temporal shifts between multivariate activity patterns evoked by the first viewing of the movie clip compared to repeated viewings (Figure 1). This model assumes that the neural response to a structured narrative stimulus consists of a sequence of distinct, stable activity patterns that correspond to event structure in the narrative (Baldassano et al., 2018). We could then identify, on a timepoint-by-timepoint basis, the extent to which viewers were activating event-specific neural patterns earlier in subsequent viewings of the movie, by drawing on their prior experience. We also compared the neural event boundaries identified by the model to event annotations in a separate group of participants. This allowed us to test whether the timing of event shifts in the brain

before and/or after repeated viewings aligned with subjective event segmentation of the narrative. Together, this approach allowed us to characterize the nature of hierarchical prediction in the brain and link it to behavioral measures of event perception.

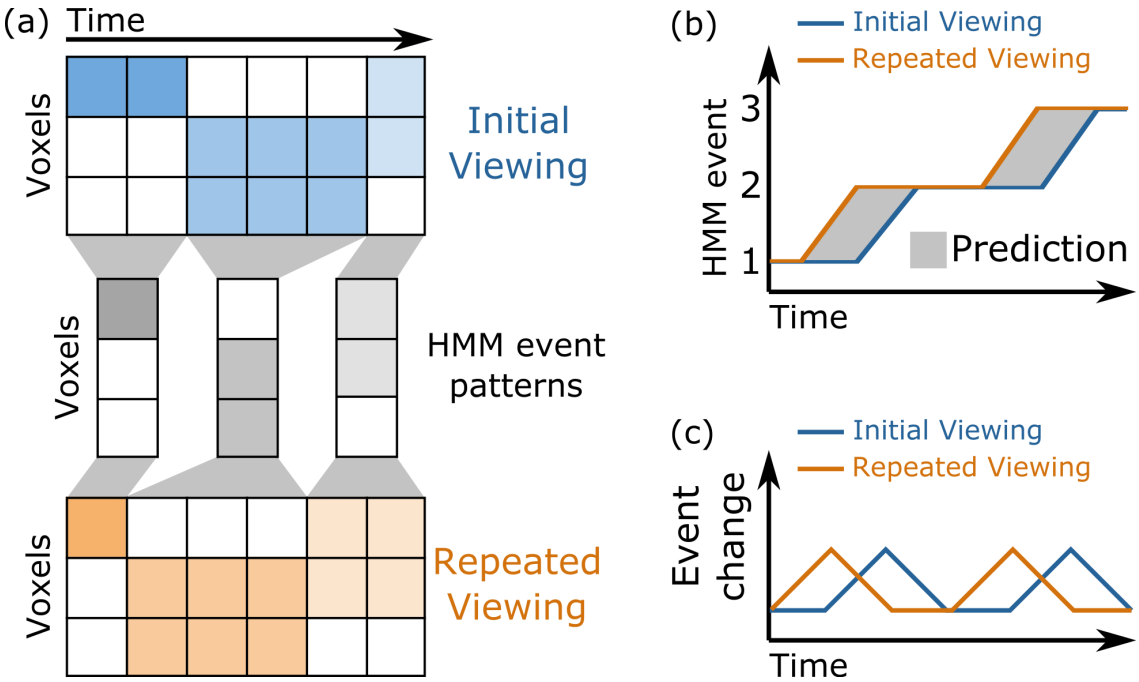

Figure 1: **Computing varying timescales of prediction by examining temporal shifts in events across multiple viewings of a movie**. (a) Using a Hidden Markov Model (HMM), fMRI activity patterns are segmented into stable events that are shared between initial and repeated viewings of a movie. The model assumes that these events have the same spatial pattern of voxel activity in both conditions, but allows the timing of event transitions to differ. (b) Regions with predictive representations are those in which event transitions occur earlier in time for repeated viewings of a stimulus compared to the initial viewing. The magnitude of this shift indicates the amount of temporal prediction. (c) Taking the temporal derivative of the event time-course plot in (b) produces a measure of the strength of event shifts at each moment in time, allowing for comparison with event boundary annotations from human observers.

3 Results

Timescales of prediction in the brain

To identify predictive signals in the brain, we examined TR-by-TR brain activity patterns during the first viewing of the movie clip and during subsequent viewings. For each spherical searchlight within the brain volume, we fit a Hidden Markov Model (HMM) jointly to both viewing conditions (initial and repeated) to identify the timing of spatial pattern changes in both cases. This model allowed us to estimate the extent to which event transitions shifted earlier in time on repeated viewings, in each region of the brain. Our analysis revealed temporal shifts in event patterns in many brain regions, including visual cortex, parahippocampal cortex, angular gyrus, retrosplenial cortex, insular cortex, temporal pole, and medial and lateral prefrontal cortex. The magnitude of this shift varied along a posterior to anterior temporal hierarchy, with the most anterior regions in the temporal pole and prefrontal cortex showing shifts of up to 12 seconds on subsequent viewings compared to the first viewing. Sensory-processing regions such as the lateral occipital cortex predicted about 3 seconds into the future, whereas the angular gyrus and ventral prefrontal cortex predict 5 and 9 seconds, respectively, into the future (Figure 2).

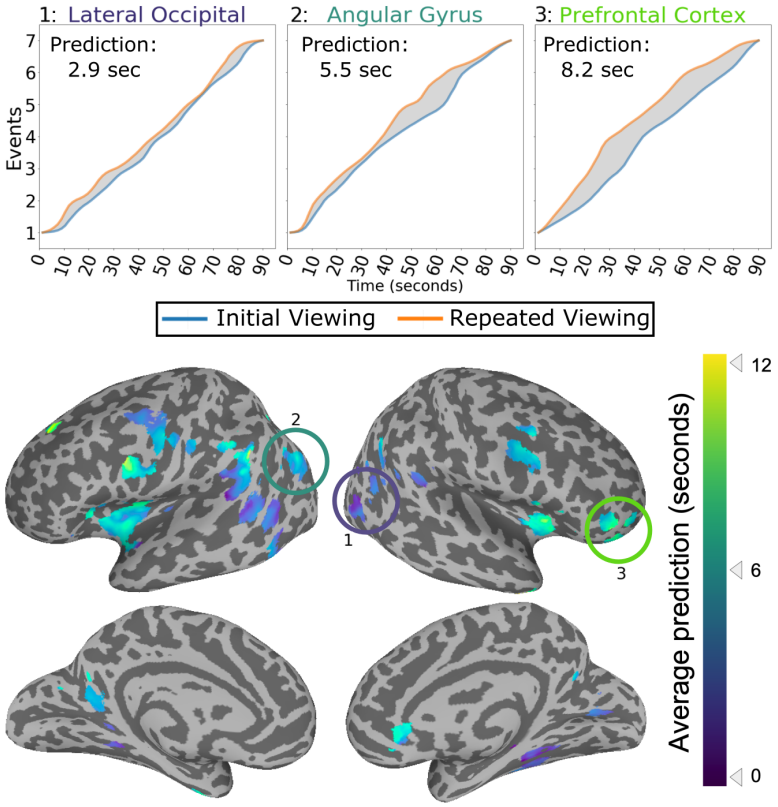

Figure 2: **Timescales of prediction vary from across the cortical hierarchy**. (Top) Multiple regions exhibited shifts in event timing between initial and repeated viewings, with event transitions shifting earlier in time with subsequent viewings (indicated by an upward shift in these event vs. time plots). (Bottom) Across the brain, prediction timescales varied from a few seconds to more than ten seconds, with the longest-timescale predictions in insular and prefrontal cortex. Statistical thresholding was conducted via bootstrapping, with correction for False Discovery Rate (FDR), $q < 0.05$

Relationship with human-annotated events

Our data-driven method for identifying event structure in fMRI data does not make use of information about the content of the stimulus, leaving open the question of how the HMM-identified event boundaries correspond to subjective event transitions in the movie. Here, we test two alternative hypotheses. First, neural event boundaries could *start* well-aligned with event boundaries in the movie and then shift earlier (indicating prediction of upcoming stimulus content). Alternatively, they may initially lag behind stimulus boundaries (reflecting a delayed response time on initial viewing) and then shift to become *better* aligned with movie scene transitions on repeated viewings.

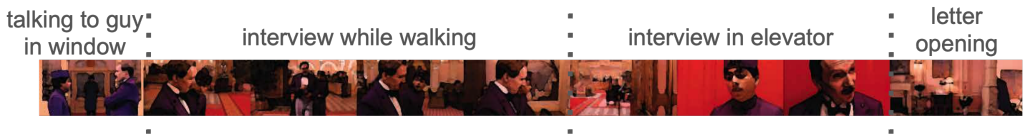

Figure 3: **An example of event annotations from *The Grand Budapest Hotel*.** Dotted lines demarcate events and phrases between the lines are brief titles given by one participant to describe each event. (Frames in this figure have been blurred to comply with copyright restrictions, but all participants were presented with the original unblurred version.)

We asked human raters to identify event transitions in the stimulus, labeling each "meaningful segment" of activity (Figure 3). To generate a hypothesis about the strength and timing of event shifts in the fMRI data, we convolved the distribution of boundary annotations with a standard Hemodynamic Response Function (HRF) as shown in Figure 4. We then explored alignment between these human-annotated event boundaries and the event boundaries extracted from the brain. For each significant cluster from the main analysis (exhibiting earlier event transitions with repeated viewing), we correlated the timecourse of HMM event transitions (see Figure 1c) with the HRF-convolved boundary annotations. Using a bootstrap analysis, we tested whether these correlations were significantly positive for initial viewing and/or repeated viewing, and whether there was a significant shift in correlation between these conditions. Two regions exhibiting significant differences in initial vs. repeated correlations are a region in lateral occipital cortex and a region in dorsal prefrontal cortex (Figure 5). For both of these regions the initial viewing data exhibits transitions near the annotated boundaries, while event transitions in repeated viewing data occur substantially earlier than stimulus transitions. A summary of results for all clusters is shown in Figure 6.

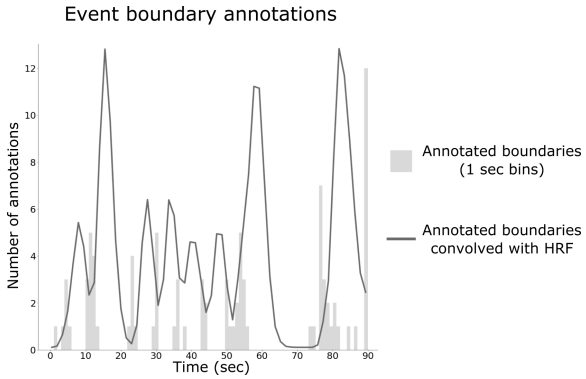

Figure 4: **Construction of behavioral and neural boundary timecoures.** The number of boundary annotations at each second of the movie clip (in gray) was convolved with an HRF to produce a continuous measure of boundary strength (black line).

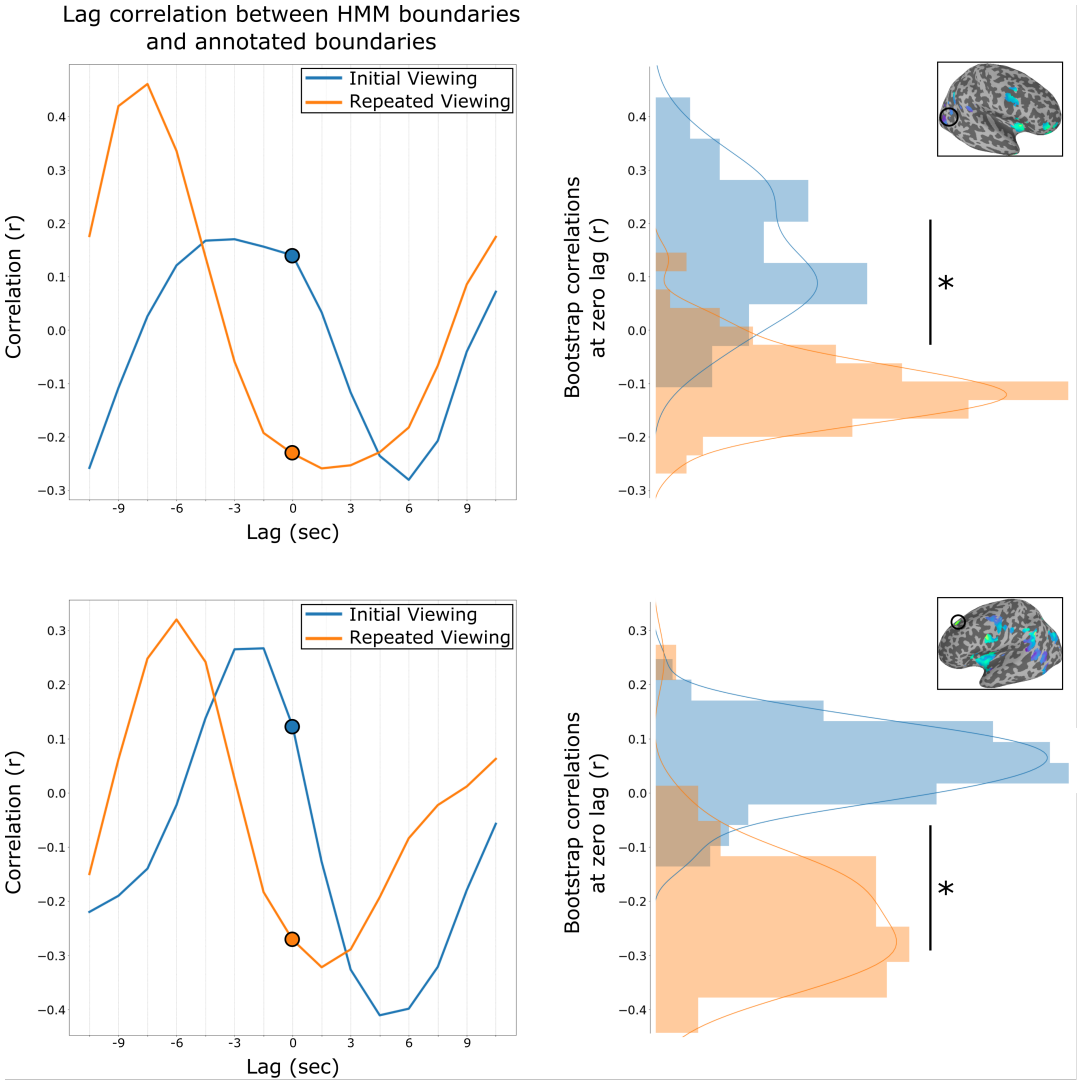

Figure 5: **Correlations between neurally-identified event transitions and human-annotated event transitions.** Lag correlation plots show the correlation between neural event boundaries and annotated event boundaries as a function of how much the timecourses are shifted with respect to one another. The correlation at 0 lag indicates the similarity between the neural and annotated event boundaries when the timecourses are aligned. Negative lags show the correlations between the timecourses when the annotated event timecourse is shifted backward in time relative to the neural event timecourse; positive lags show the correlation when the annotated event timecourse is shifted forward in time. Peaks in the lag correlation plot indicate the lag for which the correlation between the neural and annotated events is highest. At lag 0, event transitions from HMMs fit to lateral occipital (top) and dorsal prefrontal (bottom) activity are significantly more correlated to event shifts in the movie for the first viewing compared to repeated viewing. Measuring correlations at a range of lags indicates that transitions of neural events during repeated viewing occur 6-7 seconds before event transitions in the stimulus.

Clust	X	Y	Z	Vol (mm^3)	Pred (s)	r_{inital}	$r_{repeated}$	r diff
1	24.3	-80.9	11.7	159	2.6	-0.09	-0.06	-0.03
2	41.0	-66.6	10.2	324	2.9	0.14	-0.23	**0.37***
3	27.2	-42.9	-20.5	375	4.0	**0.34***	**0.38***	-0.04
4	-55.4	-45.3	8.1	925	4.0	**0.25***	-0.03	0.28
5	53.7	-57.4	28.6	30	4.8	**0.36***	0.20	0.16
6	-11.2	-54.8	11.1	61	5.2	-0.09	-0.15	0.06
7	43.7	5.6	31.1	105	5.4	0.18	**0.26***	-0.08
8	-49.4	-69.0	29.2	543	5.5	**0.31***	**0.28***	0.02
9	-42.5	1.1	12.0	1198	6.5	0.01	0.00	0.01
10	-30.3	22.5	12.0	47	7.5	-0.06	-0.16	0.10
11	30.7	43.8	-11.3	773	8.2	0.11	0.07	0.04
12	38.7	13.0	-4.1	398	9.7	-0.03	0.14	-0.17
13	-34.0	-5.2	-39.3	71	10.4	0.20	0.33	-0.13
14	-14.8	35.1	38.6	37	10.4	0.12	-0.27	**0.39***
15	42.8	4.0	-44.9	47	12.6	**0.17***	0.26	-0.09

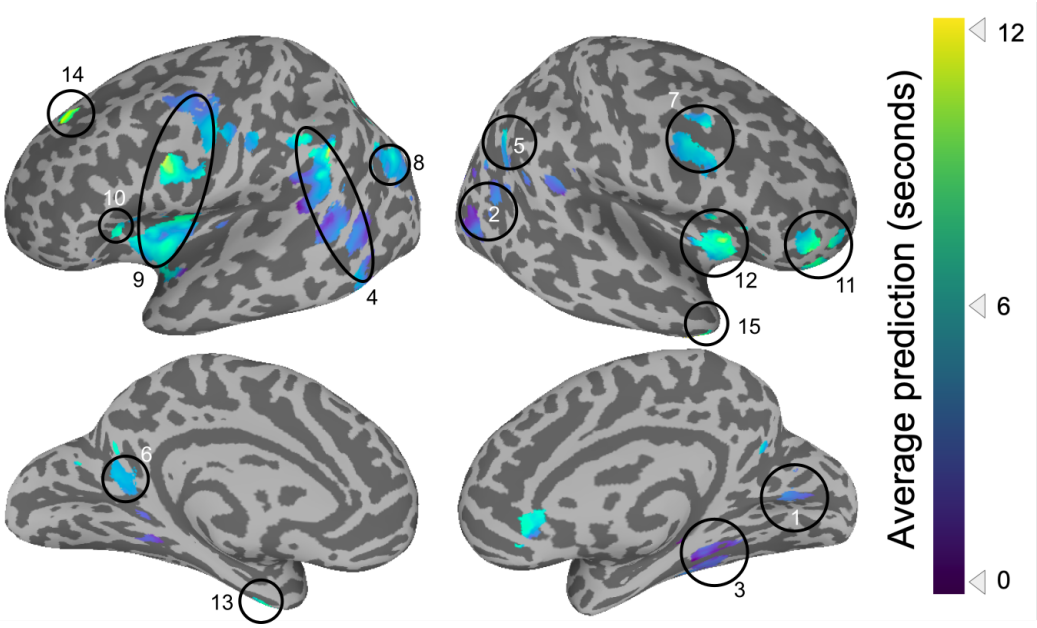

Figure 6: **Statistics for significant clusters**. For each significant cluster of at least 30 mm^3 in the main prediction analysis, we computed the center of mass (MNI), volume, average prediction timescale within cluster voxels, and correlations with annotated boundaries for initial and repeated viewing (and the difference in correlations). * $p < 0.05$.

4 Discussion

We investigated whether the brain contains a hierarchy of prediction timescales during passive viewing of a naturalistic movie. We found that regions throughout the brain exhibit flexible predictions of upcoming events in audiovisual stimuli, with activity patterns shifting earlier in time as participants repeatedly watched the same movie clip. This prediction occurred at varying timescales along the cortical hierarchy. Higher-order regions predicted further into the future than lower-order regions, and synchronization between event representations of a subset of these regions and human-annotated events shifted with learning.

Previous work has also identified cumulatively longer timescales up the cortical hierarchy, but has primarily focused on representations of the past. Lerner, Honey, Silbert, and Hasson (2011) demonstrated this effect in participants who listened to variants of a 7-minute narrative that was scrambled at different timescales (e.g., paragraphs, sentences, or words). Response reliability, measured as the correlation in BOLD activity timecourses across individuals, varied based on the timescale of scrambling, with higher-level brain regions responding consistently to only the more-intact narrative conditions. This led to the idea that higher-order brain regions contain larger "temporal receptive windows" than lower-order areas, in that their activity at a given moment is influenced by relatively more of the past. Likewise, using intracranial EEG, Honey et al. (2012) observed progressively longer temporal receptive windows in successive stages of the cortical hierarchy in participants who watched intact and scrambled versions of the movie *Dog Day Afternoon*. These findings can be described by the process memory framework (Hasson et al., 2015), where hierarchical memory timescales process, represent, and support longer and longer units of information. We found that this hierarchy also exists in the prospective direction, with the degree of predictive temporal shift increasing from posterior to anterior regions of the brain. These results extend the process memory framework, suggesting that the timescales in these regions are relevant not only for online processing and memory, but also for future prediction or simulation.

Although prior work has uncovered predictive coding in the brain, most studies have examined fixed, shorter timescales of prediction. Moreover, these shorter timescales have been studied using simple, non-narrative stimuli such as objects moving across the screen, short visual sequences, and visual pattern completion tasks (Alink et al., 2010; Ekman et al., 2017; Gavornik & Bear, 2014; Hindy et al., 2016; Kok et al., 2012, 2014). Research investigating longer timescales of prediction, such as learning future state representations in a maze task, examined single timescales up to 30 seconds ahead in OFC-VMPFC regions (Elliott Wimmer & Büchel, 2019). Studies that use narrative stimuli have examined specially-constructed texts in order to manipulate predictions about upcoming sentences; for example, Kandylaki et al. (2016) demonstrated that predictive processing of referents in narratives can be modulated by voice (passive vs active) and causality (high vs low). Our results show that in a naturalistic setting, in which structure exists at many timescales, predictions at multiple levels can occur in parallel across different brain regions. We found anticipation up to approximately 12 seconds into the future with our 90-second stimulus, but future work with stimuli of longer duration could uncover even longer timescales of prediction, on the scale of minutes. Simultaneously maintaining predictions at varying timescales could allow for flexible behaviors, because different timescales of prediction may be helpful for a variety of tasks and actions. Taking action to avoid immediate harm or danger would require shorter timescales of prediction, whereas cultivating social relationships demands predictions on longer timescales.

These results are consistent with those of Baldassano et al. (2017), in which some participants listening to an audio narrative had advance knowledge of the high-level events of the story (because they had previously watched a movie version of the narrative). Using a similar HMM approach as in this paper, the authors observed shifts in event boundaries in higher-level regions including angular gyrus, posterior medial cortex, and medial prefrontal cortex. In the current study, however, partic-

ipants were repeatedly exposed to an identical movie stimulus, allowing them to make predictions at a broad range of timescales, including the timescales of fast-changing low-level visual features. This novel approach allowed us to observe for the first time that anticipation occurs in both low- and high-level regions, with shorter-timescale predictions in visual occipital regions and the longest predictions in prefrontal cortex.

We did not observe reliable predictions in primary sensory areas, although prior studies have been able to observe predictive coding in early regions such as V1 (Alink et al., 2010; Ekman et al., 2017; Gavornik & Bear, 2014; Hindy et al., 2016; Kok et al., 2012, 2014). One possibility is that the rich, ongoing sensory input dominated relatively small prediction signals in these regions. Alternatively, such regions may predict at a timescale that is too fast to observe with fMRI. Future studies employing methods with higher temporal resolution, e.g., intracranial EEG, will be important for adjudicating between these possibilities.

Hierarchies of predictive coding could be indicative of bottom-up construction of predictions, in which each level of a predictive hierarchy passes up error signals to higher levels based on its current scale of prediction. Alternatively, as Kiebel, Daunizeau, and Friston (2008) have proposed, sensory prediction error may be guided by top-down modulation, in which higher level areas code the most likely trajectory of expected lower level sensory input. Our results are more consistent with this top-down modulation view, in which prior episodic memories can be used to pre-activate event representations throughout regions across the timescale hierarchy, rather than a prediction error account in which event representations are no longer present for expected stimulus inputs.

Our model detects anticipation as temporal shifts in events, without any change to the underlying event patterns. This view of prediction is complementary to other theories of predictive representations, in which event patterns themselves should change over time to incorporate future information. One example is the "successor representation" model from the field of reinforcement learning, which describes a representation in which each state (here, event representation) comes to include features of future events, weighted by their likelihood of occurring and their distance into the future (Dayan, 1993). Successor representations can also be constructed at multiple scales (by changing the relative weighting of events near vs far in the future). Such multi-scale representations are useful for goal-directed prediction that require multiple stages of planning (Momennejad & Howard, 2018; Brunec & Momennejad, 2019). Future work could explore how these two different theories could be integrated to model both mixing of event patterns as well as temporal shifts in the activation of these event patterns.

The detection of varying timescales of prediction in the brain can be applied to multiple domains and modalities of memory research. Future work could explore even shorter timescales using neuroimaging modalities such as EEG, or longer timescales using longer movies or narratives from TV series that span multiple episodes. Furthermore, the impact of top-down goals on the hierarchy of prediction timescales could be explored by using different tasks that require different levels of prediction, such as anticipating camera angle changes versus location changes.Our results and analysis approach provide a new framework for studying how predictive signals are distributed throughout the cortex, modulated by prior memory, and adaptive for improving comprehension and behavior.

5 Methods and Materials

5.1 Grand Budapest Hotel dataset

We used data collected by Aly et al. (2018). Thirty individuals (12 men, age: M = 23.0 years, SD = 4.2; education: M = 15.3 years, SD = 3.2; all right-handed) watched movie clips from *The Grand Budapest Hotel* while undergoing fMRI. We analyzed data from the *Intact* condition, during which participants watched a continuous 90-second clip from the movie in its original temporal order. This clip was watched 6 times, interspersed with other video clips that are not considered here. This *Intact* clip depicts an interview scene between the protagonist and his future employer inside of the Grand Budapest Hotel. Stimuli and data are available on OpenNeuro: https://openneuro.org/datasets/ds001545/versions/1.1.1

Data were acquired on a 3T Siemens Prisma scanner with a 64-channel head/neck coil using a multiband echo planar imaging (EPI) sequence (repetition time = 1.5 sec; echo time = 39 msec; flip angle = 50°; acceleration factor = 4; shift = 3; voxel size = 2.0-mm iso). T1-weighted structural images (whole-brain high-resolution; 1.0-mm iso) were acquired with an MPRAGE sequence. Field maps (40 oblique axial slices; 3-mm iso) were collected to aid registration. The fMRI scan took place over 3 experimental runs, each of which contained 2 presentations of the *Intact* movie clip (as well as other movie clips not considered here).

The first 3 EPI volumes of each run were discarded to allow for T1 equilibration. Data preprocessing was carried out in FSL, and included brain extraction, motion correction, high-pass filtering (max period = 140 sec), spatial smoothing (3-mm FWHM Gaussian kernel), and registration to standard Montreal Neurological Institute (MNI) space. After preprocessing, the functional images for each run were divided into volumes that corresponded to each of the video clips presented within that run, and only the 2 *Intact* clips within each run are considered further. Finally, each voxel's timecourse was z-scored to have zero mean and unit variance.

5.2 Event annotations by human observers

Fourteen individuals (9 men) were asked to mark event boundaries corresponding to the same 90-second *Intact* clip from *The Grand Budapest Hotel* as shown to the fMRI participants. Each participant was asked to pause the clip at the end of a meaningful segment, and to record the time and a brief title corresponding to the segment (Figure 3). Specifically, they were given the following instructions: *"The movie clip can be divided into meaningful segments. Record the times denoting when you feel like a meaningful segment has ended. Pause the clip at the end of the segment, write down the time in the spreadsheet, and provide a short, descriptive title. Try to record segments with as few viewings of the movie clip as possible; afterward, record the number of times you viewed the clip."* Although participants were allowed to watch the clip multiple times, they were instructed to minimize and report the number of viewings needed to complete the task. No participant reported watching the clip more than three times.

5.3 Detecting prediction using an Event Segmentation Model

Group-averaged fMRI data were fit with the Event Segmentation Model described by Baldassano et al., 2017. This Hidden Markov Model (HMM) assumes that 1) events are a sequence of discrete states, 2) each event is represented in the brain by a unique spatial activity pattern, and 3) all viewings of the movie evoke the same sequence of activity patterns in the same order (though possibly with different timings). We divided the six movie viewings to create a condition in which participants had no episodic memory for the stimulus (viewing 1) and a second condition in which participants had episodic memory (averaging viewings 2-6), and fit the HMM jointly to these two

datasets. This fitting procedure involved simultaneously estimating a sequence of event activity patterns that were shared between the first and repeated viewings, and estimating the probability of belonging to each of these events for every timepoint in both datasets. The model was fit with 7 events; this number was chosen to match the approximate timescale of the semantic events in the narrative, matching the mean number of events annotated by human observers (mean = 6.5).

After fitting the HMM, we obtain an event by time-point matrix for each condition (first and repeated viewings), giving the probability that each timepoint belongs to each event. We took the expectation over events at each timepoint, yielding curves (see Fig 2) showing the average event label at each timepoint, for both of the datasets. In order to compute shifts in time between first and last viewings (prediction), the area under the curve (AUC) was computed for both datasets, and we computed the amount of prediction as the AUC for repeated viewing minus the AUC for the first viewing. We then performed a one-tailed statistical test (described below) to determine whether this difference was significantly positive, indicating earlier event transitions with repeated viewing. Not only does this approach provide a way of quantifying prediction, it gives us a trajectory of the most likely event at any given time point, as well as the onset and duration of each event.

We obtained whole-brain results, and accompanying statistical analysis, using a searchlight analysis and bootstrapping. We fit the HMM to each spherical searchlight (radius = 5 voxels; stride = 5 voxels) and computed the amount of prediction using the AUC procedure described above. The searchlight analysis was repeated for each of the 100 samples comprising our bootstrap distribution. We generated bootstrap samples by randomly resampling participants, with replacement, from the dataset in the main analysis. We fit a Normal distribution to the AUC differences in these bootstrap samples, and computed a p value as the area of this Normal distribution that fell below zero. (i.e. showed later event transitions with repeated viewing, in opposition to our hypothesis). Searchlights were determined significant ($q < 0.05$) after applying FDR (False Discovery Rate) correction.

5.4 Comparison of event boundaries in brain regions to annotations

We evaluated alignment of event boundaries in the annotations and neural data by comparing the extent to which each timepoint served as an event boundary in both kinds of data. To obtain an event boundary timecourse from the annotations, we convolved the number of annotations (across all raters) at each second with the hemodynamic response function (HRF) (Figure 4). This yields a prediction about when (and how strongly) event boundaries would be present in BOLD signals. We derived a neural measure of boundary strength by re-fitting the HMM model to clusters of significant voxels in the main analysis. We took a derivative of the expected event label (Figure 1b) from the HMM results, separately for each ROI and condition (first viewing, repeated viewings). A transition between two events would be marked by a peak; stability within an event would be associated with minimal amounts of change over time. Thus, this approach provided us with a continuous measure of event change (Figure 1c). The behavioral and neural event boundary timeseries were compared using Pearson correlation to identify regions in which the degree of pattern change (measured via the HMM) was related to the boundary annotations. We created a bootstrap distribution of the correlations for the first and repeated viewings (and the difference between the two) via the same bootstrap resampling procedure used in the main analysis, and computed a p value for each correlation as the fraction of null values less than 0. We additionally performed a lag analysis, shifting the relative timing of the behavioral and neural boundary timecourses before correlating, in order to identify neural boundary signals that were shifted relative to the stimulus boundaries.

References

Alink, A., Schwiedrzik, C. M., Kohler, A., Singer, W., & Muckli, L. (2010). Stimulus predictability reduces responses in primary visual cortex. Journal of Neuroscience, 30(8), 2960–2966.

Aly, M., Chen, J., Turk-Browne, N. B., & Hasson, U. (2018). Learning naturalistic temporal structure in the posterior medial network. J Cogn Neurosci, 30(9), 1345-1365.

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.

Baldassano, C., Hasson, U., & Norman, K. A. (2018). Representation of real-world event schemas during narrative perception. , 38(45), 9689–9699. doi: 10.1523/JNEUROSCI.0251-18.2018

Brunec, I. K., & Momennejad, I. (2019). Predictive representations in hippocampal and prefrontal hierarchies. bioRxiv.

Clark, A. (2013, 05). Whatever next? predictive brains, situated agents, and the future of cognitive science. The Behavioral and brain sciences, 36, 1-24. doi: 10.1017/S0140525X12000477

Dayan, P. (1993). Improving generalization for temporal difference learning: The successor representation. Neural Computation, 5(4), 613-624.

de Lange, F. P., Heilbron, M., & Kok, P. (2018, 2020/06/22). How do expectations shape perception? Trends in Cognitive Sciences, 22(9), 764–779. doi: 10.1016/j.tics.2018.06.002

den Ouden, H. E. M., Daunizeau, J., Roiser, J., Friston, K. J., & Stephan, K. E. (2010). Striatal prediction error modulates cortical coupling. Journal of Neuroscience, 30(9), 3210–3219. doi: 10.1523/JNEUROSCI.4458-09.2010

Ekman, M., Kok, P., & de Lange, F. P. (2017). Time-compressed preplay of anticipated events in human primary visual cortex. Nature Communications, 8(1), 15276. doi: 10.1038/ncomms15276

Elliott Wimmer, G., & Büchel, C. (2019). Learning of distant state predictions by the orbitofrontal cortex in humans. Nature Communications, 10(1), 2554.

Friston, K. (2005). A theory of cortical responses. Philosophical transactions of the Royal Society of London. Series B, Biological sciences, 360, 815-36. doi: 10.1098/rstb.2005.1622

Gavornik, J. P., & Bear, M. F. (2014). Learned spatiotemporal sequence recognition and prediction in primary visual cortex. Nature Neuroscience, 17(5), 732–737.

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. doi: https://doi.org/10.1016/j.tics.2015.04.006

Hindy, N., Ng, F., & Turk-Browne, N. (2016, 04). Linking pattern completion in the hippocampus to predictive coding in visual cortex. Nature neuroscience, 19. doi: 10.1038/nn.4284

Honey, C., Thesen, T., Donner, T., Silbert, L., Carlson, C., Devinsky, O., ... Hasson, U. (2012, November 8). Slow cortical dynamics and the accumulation of information over long timescales. Neuron, 76(3). doi: 10.1016/j.neuron.2012.10.024

Kandylaki, K. D., Nagels, A., Tune, S., Kircher, T., Wiese, R., Schlesewsky, M., & Bornkessel-Schlesewsky, I. (2016). Predicting "when" in discourse engages the human dorsal auditory stream: An fmri study using naturalistic stories. Journal of Neuroscience, 36(48), 12180–12191.

Kok, P., Failing, M., & de Lange, F. (2014, 01). Prior expectations evoke stimulus templates in the primary visual cortex. Journal of cognitive neuroscience, 26. doi: 10.1162/jocn$_{a0}$0562

Kok, P., Jehee, J. F., & de Lange, F. P. (2012). Less is more: Expectation sharpens representations in the primary visual cortex. Neuron, 75(2), 265 - 270.

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.

Momennejad, I., & Howard, M. W. (2018). Predicting the future with multi-scale successor representations. bioRxiv.

Summerfield, C., & Egner, T. (2009, 09). Expectation (and attention) in visual cognition. Trends in cognitive sciences, 13, 403-9. doi: 10.1016/j.tics.2009.06.003