

2011

Prediction error associated with the perceptual segmentation of naturalistic events

Jeffrey M. Zacks
Washington University School of Medicine in St. Louis

Christopher A. Kurby
Washington University in St Louis

Michelle L. Eisenberg
Washington University in St Louis

Nayiri Haroutunian
Washington University in St Louis

Follow this and additional works at: https://digitalcommons.wustl.edu/open_access_pubs

Please let us know how this document benefits you.

Recommended Citation

Zacks, Jeffrey M.; Kurby, Christopher A.; Eisenberg, Michelle L.; and Haroutunian, Nayiri, "Prediction error associated with the perceptual segmentation of naturalistic events." *Journal of Cognitive Neuroscience*. 23, 12. 4057-4066. (2011).

https://digitalcommons.wustl.edu/open_access_pubs/3356

This Open Access Publication is brought to you for free and open access by Digital Commons@Becker. It has been accepted for inclusion in Open Access Publications by an authorized administrator of Digital Commons@Becker. For more information, please contact vanam@wustl.edu.

Prediction Error Associated with the Perceptual Segmentation of Naturalistic Events

Jeffrey M. Zacks¹, Christopher A. Kurby^{1,2}, Michelle L. Eisenberg¹,
and Nayiri Haroutunian¹

Abstract

■ Predicting the near future is important for survival and plays a central role in theories of perception, language processing, and learning. Prediction failures may be particularly important for initiating the updating of perceptual and memory systems and, thus, for the subjective experience of events. Here, we asked observers to make predictions about what would happen 5 sec later in a movie of an everyday activity. Those points

where prediction was more difficult corresponded with subjective boundaries in the stream of experience. At points of unpredictability, midbrain and striatal regions associated with the phasic release of the neurotransmitter dopamine transiently increased in activity. This activity could provide a global updating signal, cuing other brain systems that a significant new event has begun. ■

INTRODUCTION

Humans and other species depend heavily on predictions about the near future to guide behavior. Predictions that allow one to anticipate features such as the movements of objects and the behaviors of other animals are of great adaptive benefit: They allow an organism to anticipate threats and opportunities and to respond appropriately before it is too late. In perception, predictions about forthcoming visual content guide eye movements, attention, and the neural processing of objects' identity and location (Bar, 2009; Enns & Lleras, 2008; Summerfield et al., 2006). In language comprehension, predictions about likely words guide word recognition and syntactic parsing (Elman, 2009; Pickering & Garrod, 2007). When learning to perform a task, predictions about the reward value of potential actions are used to improve performance (Maia, 2009).

Errors in prediction are a valuable source of information about whether an organism's representation of the environment is effective. Research on *reinforcement learning* has shown that by monitoring errors in prediction while behaving, an organism can form an accurate model of the environment that supports adaptive behavior (Maia, 2009; Niv & Schoenbaum, 2008). When representations fit events in the world well, prediction error should be low; when they fit poorly, prediction error should rise. One intriguing possibility is that transient increases in prediction error could be used to regulate attentional control over time. According to one recent model (Zacks, Speer, Swallow, Braver, & Reynolds, 2007), working memory representations of the current event guide perceptual predic-

tions about the immediate future. These predictions are checked against what happens next in the perceptual stream; most of the time, perceptual predictions about what happens next are accurate. From time to time, however, activity becomes less predictable, causing a spike in prediction errors. These spikes in prediction error are fed back to update working memory and reorient the organism to salient new features in the environment. According to this model, the increase in prediction error and consequent updating results in the subjective experience of an event boundary in perceptual experience. Computational simulations support the viability of this mechanism: Recurrent neural networks were trained to perform one-step prediction tasks in an environment consisting of recurring events. Increases in prediction error were reliable predictors that a new event had begun and could be used to adaptively update memory representations, thus improving prediction performance (Reynolds, Zacks, & Braver, 2007).

The segmentation of ongoing activity into discrete events is a highly salient feature of subjective experience—and one that is tightly related to memory and cognition. In the laboratory, human subjects can easily mark off the time point at which they perceive one event to end and another to begin, and such judgments are reliable across and within observers (Speer, Swallow, & Zacks, 2003; Newtonson, 1976). Event boundaries during perception predict what will be remembered later (Zacks, Speer, Vettel, & Jacoby, 2006; Newtonson, 1976) and are associated with transient brain responses in the parietal, temporal, and frontal lobes (Zacks et al., 2001). Importantly, the brain responses are observed independently of whether observers are deliberately attending to segmentation, suggesting that these mechanisms are a normal concomitant of ongoing perception.

¹Washington University in Saint Louis, ²Grand Valley State University

One candidate system for signaling prediction error during ongoing perception is the midbrain phasic dopamine system (MDS), which includes dopamine cells in the substantia nigra (SN) and ventral tegmental area (VTA) (Schultz, 1998). The MDS projects broadly and directly to the cortex and indirectly through the striatum. Cells in the MDS show strong responses to at least three distinct sorts of unpredictability: unexpected reward, uncertainty about outcomes, and salience (Schultz, 1998, 2007; Horvitz, 2000). Responses to unexpected rewards can be used to adjust behavior to maximize reward (Maia, 2009). Uncertainty responses may potentiate the dopamine pathway, fine-tuning the prediction error signal (Schultz, 2007). Salience responses may serve as a cue to the organism to orient to new potential rewards (Kakade & Dayan, 2002) or to learn features of its environment that are not currently connected to reward but that could become important for reward later (Horvitz, 2000).

Human fMRI studies provide evidence of a role for the MDS in signaling all three kinds of uncertainty. In classical conditioning paradigms, the MDS and its striatal targets have been found to track deviations from expected reward (D'Ardenne, McClure, Nystrom, & Cohen, 2008; McClure, Berns, & Montague, 2003). In studies of classification learning tasks without an explicit reward, the MDS has been found to increase in activity when outcomes are unpredictable and when unexpected outcomes occur (Aron et al., 2004; Volz, Schubotz, & von Cramon, 2003). In target detection tasks, MDS responses to salient distractor stimuli have been observed, even when the distractors do not require any response (Zink, Pagnoni, Chappelow, Martin-Skurski, & Berns, 2006; Zink, Pagnoni, Martin, Dhamala, & Berns, 2003).

Currently, there is a debate about whether the MDS codes prediction error about reward per se or signals multiple types of prediction error and whether multiple mechanisms involving the MDS and striatum likely are involved in signaling prediction error (Gläscher, Daw, Dayan, & O'Doherty, 2010; Schultz, 2010; Horvitz, 2000; Schultz &


Dickinson, 2000). However, the various accounts are consistent with the possibility that transient increases in prediction error are signaled by the MDS to the cortex, leading to memory updating and the subjective experience that a new event has begun.

The hypothesis that subjective event boundaries correspond to transient unpredictability leads to several specific proposals: First, if observers are explicitly asked to make predictions about the near future, this should be more difficult when predicting across a subjective event boundary. Second, attempting to make predictions across event boundaries should activate the MDS. Third, the MDS should also be more active when information from a newly begun event provides a signal that one's predictions are incorrect. To test these proposals, we presented human observers with movies of naturalistic activities, measured their ability to predict what would happen in the movie in the near future, and measured activity with fMRI.

METHODS

Human observers watched movies depicting everyday events (washing a car, building a LEGO model, putting up a tent, washing clothes, and planting a window box). Three experiments probed behavioral performance, and one measured brain activity with fMRI. In all experiments, the movies were stopped approximately once per minute, and the participants were asked to make a prediction about what would be on the screen in 5 sec. The stopping points were chosen so as to occur 2.5 sec before a natural event boundary or 2.5 sec before a natural event middle based on judgments of a previous group of observers (Figure 1). After each prediction was made, the movie restarted from the pause point. This provided feedback about the participants' prediction by revealing the correct frame 5 sec after the movie restarted (Figure 2A). This paradigm was inspired by studies of probabilistic classification learning (Aron et al., 2004) and adapted for prediction during continuous activity.

Figure 1. Illustration of the selection of within-event and across-event trials. The black curve plots the probability density of segmentation for the LEGO movie. The red arrows indicate the across-event condition test points, and the blue arrows indicate the within-event condition test points.



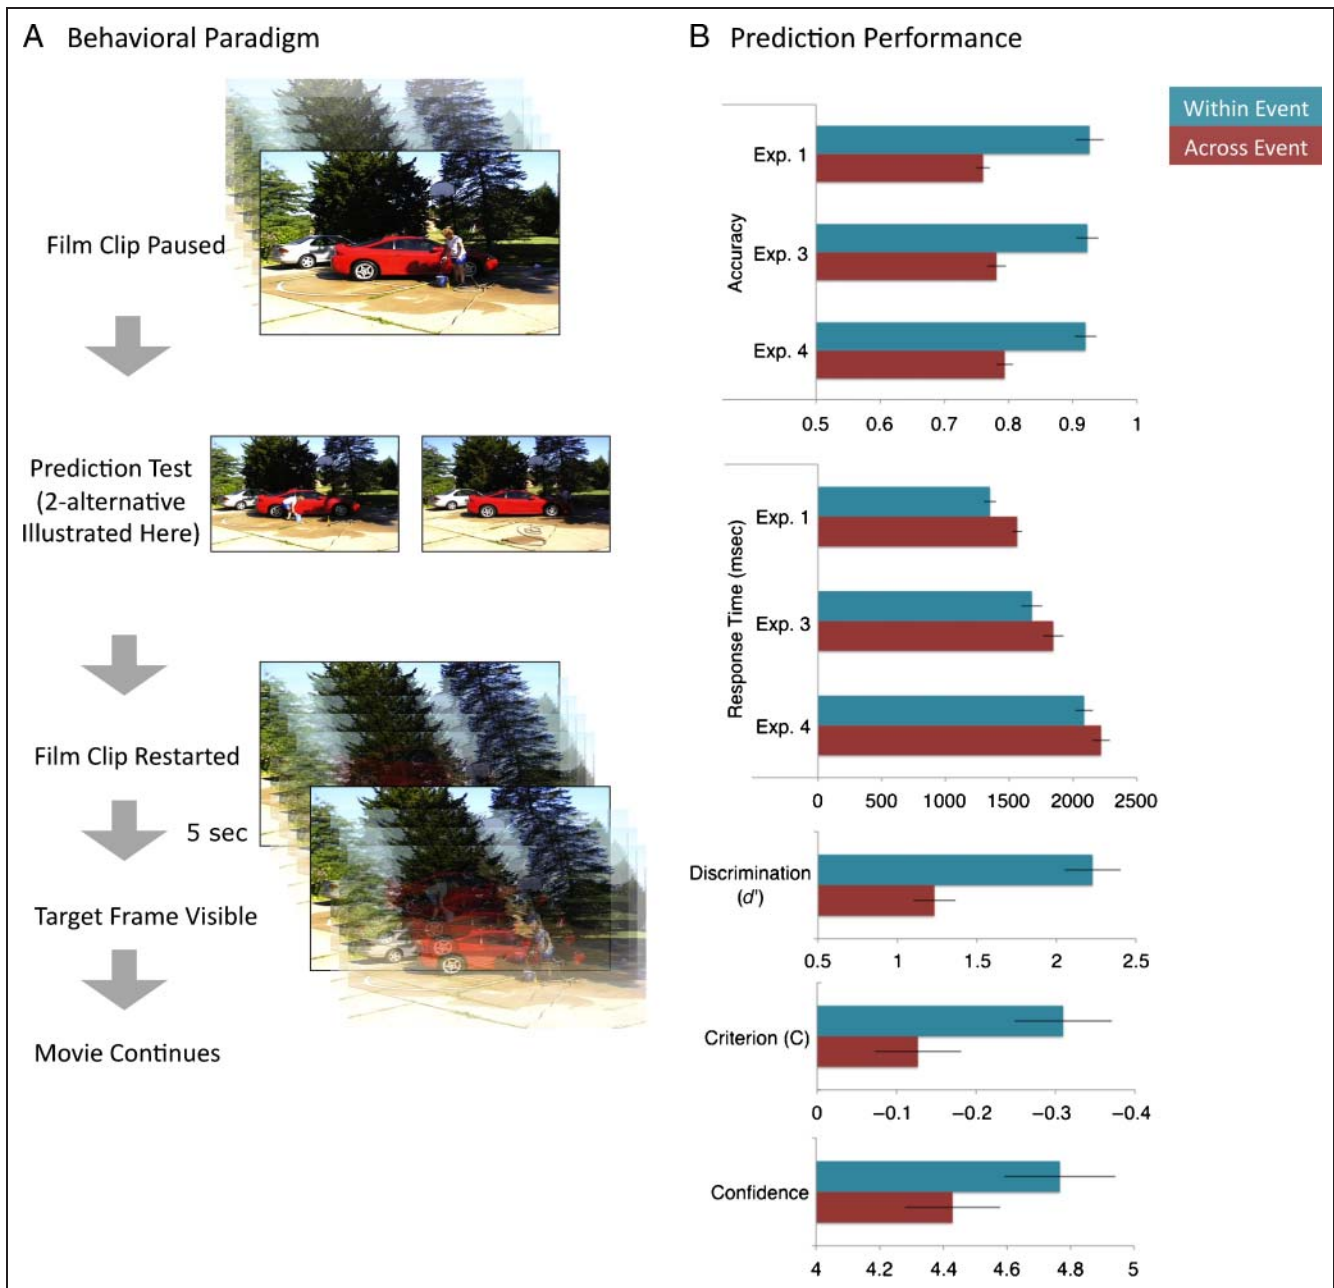


Figure 2. Predicting the near future is more difficult at event boundaries. (A) Participants viewed movies of everyday events (247- to 432-sec duration). Eight times during each movie, the movie was paused and the viewer was asked to predict what would be on the screen in 5 sec. Experiments 1 and 3 used a two-alternative forced-choice procedure, in which participants chose between the correct picture and a similar appearing picture taken from a different movie. (In this example, the picture on the left is correct.) Experiment 2 used a yes-no procedure, in which one of the two possible pictures was presented and participants were required to assess whether it was correct. In Experiment 3, participants made an assessment of their prediction confidence using a 6-point Likert-type scale before the two test pictures were presented. (B) By all measures, prediction performance was worse for across-event trials than within-event trials. Statistical tests were conducted using both participants and items as the random effect. In Experiments 1, 3, and 4, forced-choice accuracy was lower (top graph) and responses were slower (second graph) when predicting across events. [By participants: smallest $t(23) = 6.68, p < .001$; by items: smallest $t(38) = 2.84, p = .007$. For response time, by participants: smallest $t(23) = 3.28, p = .003$; by items: smallest $t(38) = 2.42, p = .007$.] Experiment 2 was analyzed using signal detection analysis (Macmillan & Creelman, 1991), which showed that discrimination between correct and incorrect pictures was lower across than within events (third graph; by participants: $t(23) = 4.82, p < .001$; by items: $t(38) = 3.56, p = .001$). Moreover, participants' decision criteria were more conservative across than within events, indicating that, overall, pictures taken from new events were judged less likely to be correct, although this effect was significant only by participants (fourth graph; by participants: $t_{23} = 2.71, p = .013$; by items: $t(38) = 0.98, p = .335, ns$). In Experiment 3, confidence judgments made before presentation of the test pictures indicated that participants were less confident in their ability to predict as they approached a new event (bottom graph; by participants: $t(23) = 6.81, p < .001$; by items: $t(38) = 2.42, p = .021$. (Error bars depict standard errors.)

Participants

Participants were drawn from the Washington University community. Participants in Experiments 1–3 received \$10 or partial credit for a course requirement. Participants in Experiment 4 received \$25/hr. Experiment 1 had 52 participants (31 women, ages = 18–54 years), Experiment 2 had 24 participants (14 women, ages = 18–25 years), Experiment 3 had 24 participants (14 women, ages = 18–22 years), and Experiment 4 had 25 participants (12 women, ages = 19–34 years).

Materials and Stimulus Presentation

Movies of everyday activities were selected from a previous study of event perception in younger and older adults (Kurby & Zacks, 2011). The five activities used were washing a car (432 sec), building a LEGO model (247 sec), putting up a tent (378 sec), washing clothes (300 sec), and planting a window box (354 sec). Each was filmed with a digital camera from a fixed head-height perspective, with no edits, zooms, or camera motion. In the previous study, neurologically healthy older (60–89 years) and younger (18–23 years) adults segmented the movies into meaningful events. They were instructed to segment while viewing by pressing a button whenever, in their judgment, one meaningful unit of activity ended and another began. For the data used in the present analyses, participants were asked to identify the largest units they found meaningful. We used these data to identify normative event boundaries. First, we estimated the probability density of segmentation throughout each movie using gaussian kernel density estimation (3-sec bandwidth). We then defined a 5-sec window around each local maximum and minimum and computed the proportion of participants that segmented the movie within this window. On the basis of these segmentation proportions, we selected the top four maxima as across-event condition test points and the bottom four minima locations as the within-event condition test points (see Figure 1). In some cases, top candidate maxima occurred at the end of the movie, likely because of the salient change in activity caused by the actor leaving the scene. We excluded these locations as candidate test points. This selection procedure resulted in a total of 20 across-event test points (mean segmentation proportion = .59, $SD = .11$) and 20 within-event test points (mean segmentation proportion = .01, $SD = .01$). The average temporal distance between test points was 39.34 sec ($SD = 27.10$ sec, min = 7.14 sec, max = 125.19 sec).

In Experiments 1–3, stimuli were presented on an LCD monitor by a Macintosh computer (Apple, Cupertino, CA) with PsyScope X software (psy.ck.sissa.it). In Experiment 4, stimuli were rear-projected onto a screen placed at the foot of the scanner bore and were viewed through a mirror attached to the scanner head coil. In all experiments, participants responded using a button box.

The task structure for Experiment 1 was as follows: Participants watched each movie from beginning to end, with interruptions for prediction trials (eight per movie). For each prediction trial, the movie was paused 2.5 sec before the local maximum (for across-event trials) or minimum (for within-event trials) in segmentation probability. The screen was cleared, and two test images appeared side by side. One image was the frame 5 sec subsequent to the pausing of the movie, and the other was a foil selected from a similar movie filmed with the same actor performing a related activity in the same setting (see Figure 2); foils were taken from the same proportional time in the alternate activity. Participants selected which of the two images they believed was about to appear in the movie. Participants were trained to make their responses within 5 sec; failure to respond in this window resulted in presentation of a timeout message. After responding, participants were given immediate feedback in the form of a text display (“RIGHT” or “WRONG”). The movie then was restarted from the pause point. The movie’s resumption provided further feedback regarding the participants’ predictions, as they then saw the frame that actually did appear. Before beginning the main task, they practiced using a movie of a woman making a sandwich (147 sec, with six prediction trials).

The task structure for Experiment 2 was identical to that of Experiment 1, except that, rather than two pictures, only one was presented, and the participants were asked to judge whether that picture would appear shortly. On half of the trials, the correct response was “yes.” (Assignment of trials to “yes” or “no” was randomized once; half of the participants received the randomized list and the other half received a list with the trial assignments reversed.) Again, immediate feedback was given, this time indicating whether the response was a correct positive, a missed positive, a correct negative, or a missed negative.

For Experiment 3, the task structure was identical to Experiment 1 with two exceptions. First, before the two alternatives were shown, participants reported their subjective confidence in their ability to predict the upcoming activity. A 6-point Likert-type scale was presented, with the ends marked 1 (*not at all*) and 6 (*extremely well*). The top of the screen presented a reminder of the instructions that read, “How well do you think you can predict what will happen in a few seconds?” Participants indicated their confidence by pressing one of six marked buttons on a button box. Then, after giving a confidence judgment, the two pictures were presented and participants made their prediction as in Experiment 1. Second, participants did not receive explicit feedback about their predictions. This was done to reduce the chance that feedback would influence later confidence judgments.

For Experiment 4, the task structure was identical to Experiment 1 with three exceptions. First, no explicit feedback was given after each picture was selected; feedback was provided solely by the resumption of the movie. Second, to allow for separation of the brain response to

the prediction trials from the response to movie resumptions, we introduced a jitter (2–10 sec) between the offset of the two picture alternatives and the restarting of the movie. Third, the two picture alternatives remained on-screen for the full 5-sec response window rather than offsetting when the participant responded.

MRI Scanning and Data Analysis

For Experiment 4, MRI scanning was conducted on a 3 T Vision scanner (Siemens, Erlangen, Germany). Functional images were collected with an echoplanar pulse sequence at a rate of one frame every 2.048 sec (slice repetition time [TR] = 64 msec, echo time [TE] = 25). To provide optimal contrast for distinguishing the SN and VTA, we acquired a proton density-weighted turbo spin echo sequence. The sequence used a double-echo procedure to also acquire a T2-weighted image (slice TR = 8.04 sec, TE = 18/105 msec). The T2-weighted image was used together with a high-resolution T1-weighted

anatomical scan for atlas normalization and visualization (slice TR = 2.1 sec, TE = 3.93 msec).

Before analysis, the functional data were processed to correct slice-to-slice timing offsets, normalize slice-to-slice intensity differences, correct for within-run and between-run motion, and normalize the whole-brain image intensity to a mode of 1000. The functional data were aligned to the T2 image, and these were aligned to the individual's high-resolution structural images. The functional and structural data were resampled to a standard stereotaxic space (Talairach & Tournoux, 1988) using 3.0-mm isotropic voxels with an atlas representative target constructed using the methodology described by Buckner and colleagues (2004) and smoothed with a Gaussian kernel (6 mm FWHM).

For each participant, the left and right SN were traced by hand on the high-resolution T1-weighted image using the proton density-weighted image as a reference (see Figure 3A). The right and left SN are hyperintense areas in the proton density image medial to the hypointense cerebral peduncles and lateral to the hypointense red

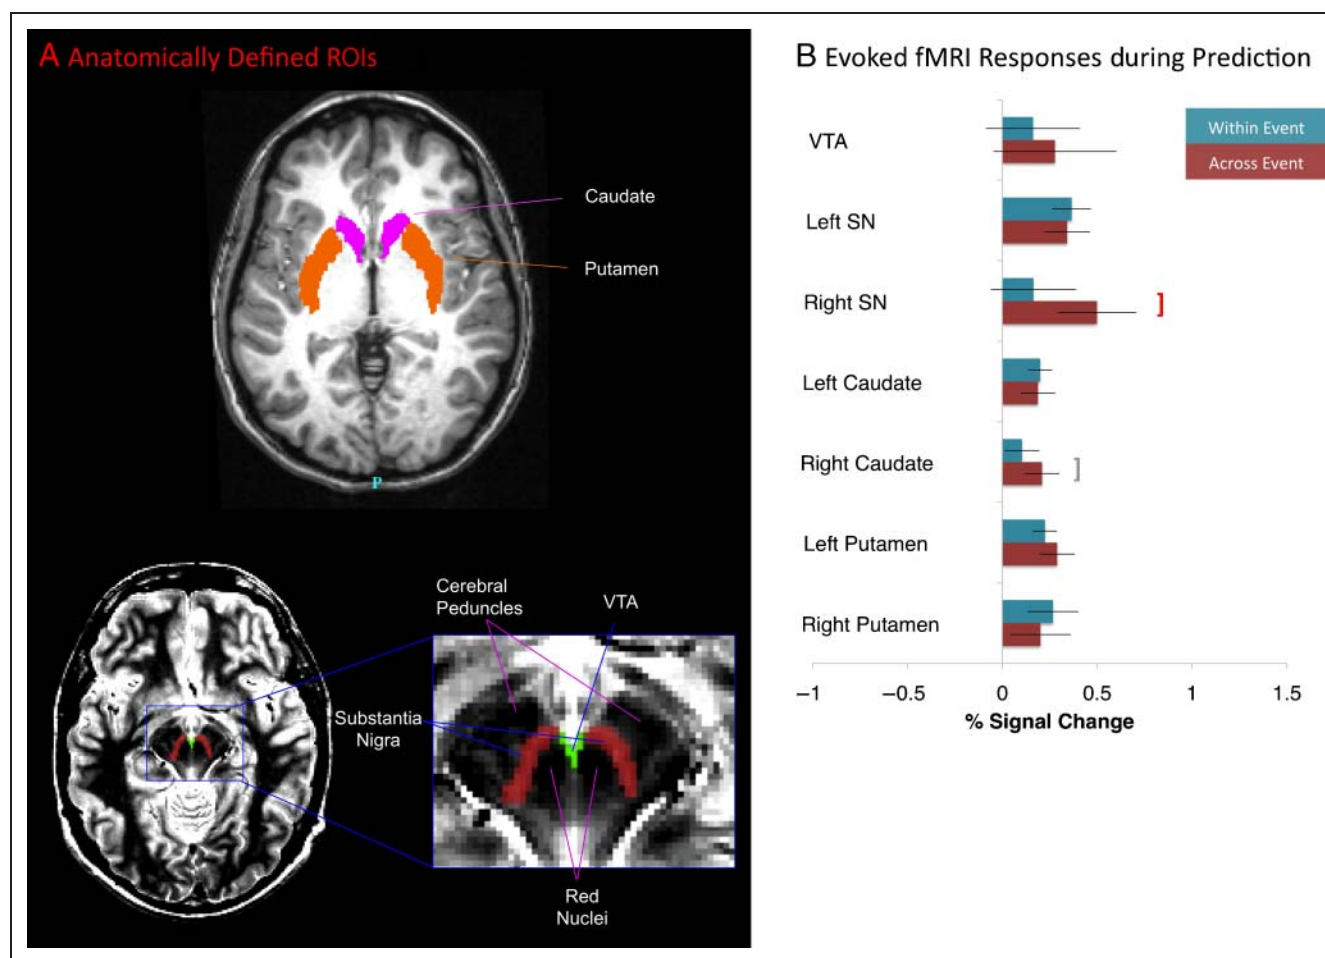


Figure 3. The MDS and its striatal targets are activated when attempting to predict features of a new event. (A) For each participant, the SN and VTAs were identified by manual tracing using Caret (Van Essen et al., 2001), and the caudate nucleus and putamen were identified using FreeSurfer (Jovicich et al., 2009; see Supplementary Methods). (B) Magnitude of evoked fMRI response when attempting to predict. Significant responses during the prediction task were observed in the left SN, caudate, and putamen. For the right SN, activity was significantly greater for prediction across than within events (red bracket); for the right caudate, this trend approached statistical significance (gray bracket). (Error bars depict standard errors.)

nuclei and hyperintense interpeduncular fossa. The SN extends along the length of the cerebral peduncles. The VTA is the hyperintense area bounded laterally by the SN and red nuclei and bordering the interpeduncular fossa (D'Ardenne et al., 2008; Haber & Gdowski, 2004; Steward, 2000). The right and left SN and the VTA were traced in each slice of the brain using the volume segmentation edit voxels function of Caret. The areas were traced in an axial view and then adjusted using coronal and sagittal views. The caudate nucleus and putamen were identified from the T1-weighted anatomic image using the automated method implemented in FreeSurfer (Jovicich et al., 2009).

fMRI data analyses were based on the general linear model (GLM), and all analyses were performed with FIDL (www.nil.wustl.edu/~fidl/). We simultaneously modeled the brain response to the prediction task and to the subsequent restarting of the movie. For each participant we constructed a GLM with effects coding for the within-event and across-event prediction trials and for the within-event and across-event restarting of the movies. Activation during the prediction task was modeled using an assumed hemodynamic response function (Boynton, Engel, Glover, & Heeger, 1996). Predictor variables for the within-event and across-event conditions were constructed by creating a train of impulses time-locked to the trials with durations equal to the participant's response time and then convolving each train of impulses with the model hemodynamic response function. This approach assumes that the interval of relevant brain activity corresponds to the period from the onset of the trial until the participant's response, and it controls for differences in response time that may contribute to differences in the estimated brain response. Activation during the subsequent restarting of the movie was modeled using a finite impulse response basis set to estimate the time course of brain activity (Ollinger, Corbetta, & Shulman, 2001). This approach was taken because it was not reasonable to specify a priori the time course of brain activity during the movie restarting and consequent confirmation or disconfirmation of predictions. Time courses of 10 frames were estimated beginning with the onset of the movie. The GLMs also included variables coding for effects of no interest, specifically scan-to-scan differences in baseline and linear trends within each scan. Error trials were excluded from the analysis because they were not frequent enough to model reliably.

For the region-wise analyses, GLMs were fit to the BOLD time course for each participant for each ROI. Analyses with participant as the random effect were carried out using t tests for the magnitude of response during the prediction trials and ANOVAs for the time course of response during the subsequent movie resumption.

For the whole-brain analyses, the GLM fitting and random effects analyses used the same form, except that the unit of analysis was the voxel rather than the region. To correct for multiple comparisons across voxels, t statistics were converted to z statistics and thresholded to include

only clusters of two or more voxels with a z of greater than 4.5. ANOVA F statistics were converted to z statistics and thresholded to include only clusters of five or more voxels with a z of greater than 4.5. These thresholds have been shown in simulation studies to control the overall probability of a false positive response at $p = .05$ (McAvoy, Ollinger, & Buckner, 2001).

RESULTS

Prediction Performance

All four experiments revealed that prediction was more difficult when predictions crossed an event boundary. In Experiment 1, participants made their judgments by selecting from two pictures, one of which was the actual frame from 5 sec later in the movie; the other picture was a foil taken from a similar movie involving the same actor and setting. Predictions were slower and less accurate when they crossed an event boundary (Figure 2B, top two graphs). In Experiment 2, participants made a prediction about whether a single picture was the frame they would see 5 sec later; half of the time, the picture was the one they would see, and half of the time, it was a foil. We used rates of correct identification and false alarms to estimate measures of prediction ability (Discrimination, d') and bias to respond "no" (Criterion, C) using signal detection theory (Macmillan & Creelman, 1991). Prediction ability was lower when predicting across events (Figure 2B, third graph). Furthermore, participants showed a stronger bias to respond "no," indicating that when predicting across an event boundary, any picture seems less likely to actually occur (Figure 2B, fourth graph). In Experiment 3, the two-alternative procedure of Experiment 1 was used but participants were also required to estimate their confidence in their predictions before each pair of test alternatives was presented. Participants reported lower confidence when trying to predict across events (Figure 2B, fifth graph). (Experiment 3 also replicated the accuracy and response time effects of Experiment 1; see Figure 2B, top two graphs.)

Midbrain Dopaminergic System Activation

To test the proposal that the MDS is engaged during an attempt to predict across an event boundary, we measured brain responses time-locked to the presentation of the two test pictures depicting possibilities for what might happen in 5 sec (see Figure 2A). The results are summarized in Figure 3B. Across-event responses were significantly greater than within-event responses in the right SN ($t_{24} = 2.2, p = .04$); this difference approached significance in the right caudate ($t_{24} = 2.0, p = .05$). For across-event trials, significant responses were observed in the left and right SN, left and right caudate, and the left putamen (smallest $t_{24} = 2.2, p = .04$). For within-event trials, significant responses were observed in the left

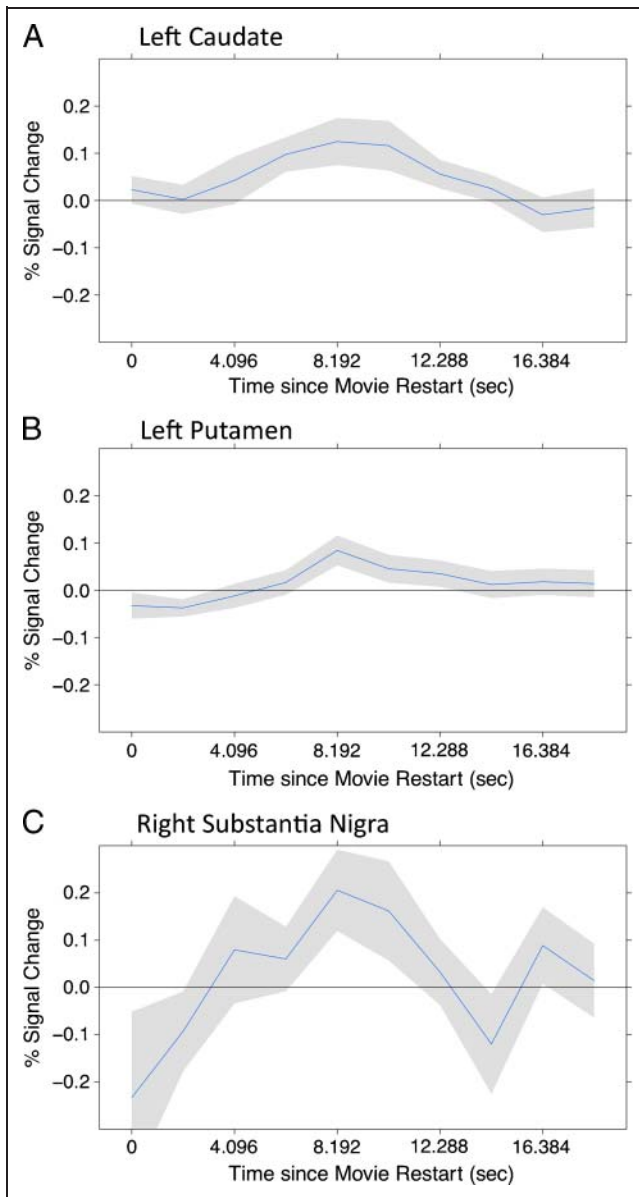


Figure 4. Disconfirming information from the beginning of a new event was accompanied by increased fMRI response in the MDS and its striatal targets. Each panel plots the difference in the fMRI response to restarting the movie in an across-event trial compared with a within-event trial. The new event's beginning was 2.5 sec after the movie onset. Between 8 and 10 sec after the movie onset, a maximal difference response is observed. (Gray regions depict standard errors.)

SN, caudate, and putamen and in the right putamen (smallest $t_{24} = 2.1, p = .05$).

To test the proposal that the MDS is engaged when information from a newly begun event provides a signal that one's predictions are incorrect, we measured brain responses time-locked to the restarting of the movie after each prediction trial. Overt errors were rare (Figure 2B, top graph), precluding a comparison between error and correct trials. However, because overt errors were more likely in the across-event condition, it is also likely that participants experienced disconfirming evidence in this

condition, even in the cases when they were uncertain or guessing but answered correctly. Therefore, we compared activity when the movie restarted in the across-event condition to that in the within-event condition.


Significant differences between across-event and within-event trials in the time course of activity during the 18.4 sec following the movie onset were observed in two regions (Figure 4): the left caudate ($F_{9, 216} = 3.02, p = .002$) and the left putamen ($F_{9, 216} = 2.73, p = .005$). This difference approached significance in the right SN ($F_{9, 216} = 1.87, p = .06$). In all three areas, activity increased during across-event trials relative to within-event trials.

Whole-brain Analyses

To investigate the responses of small brain regions, particularly the MDS ROIs, region-specific analyses are necessary. To explore the correlates of these effects across the brain, we also conducted whole-brain analyses (see MRI Scanning and Data Analysis). Responses time-locked to the presentation of the pictures were greater during across-event trials than during within-event trials in several brain regions: the juncture of the parietal, temporal, and occipital lobes bilaterally and in a small region in the right anterior temporal cortex (Figure 5). No regions showed the opposite pattern. During the intervals after restarting the movie, these regions also showed differences between within- and across-event trials, as did others bilaterally in the posterior occipital, temporal, and parietal cortex, the premotor cortex, and the left insular cortex. Similar to the responses observed in the left caudate, left putamen, and right SN, these responses were greater for across-event trials, consistent with a stronger response to disconfirming information (Figure 6). The regions identified in these analyses correspond well with those identified in previous studies of event boundaries in film (Zacks et al., 2001) and text comprehension (Whitney et al., 2009; Speer, Reynolds, & Zacks, 2007). They also include much of the dorsal and ventral attentional systems associated with attentional reorienting (Corbetta, Patel, & Shulman, 2008). The STS, which showed the strongest effects, is also strongly activated at event boundaries and has been proposed to be involved in predicting the endpoints of complex biological motion trajectories (Frith, 2007).

Could the effect observed here of event boundaries on behavioral performance and neural activity during the prediction task arise simply because of differences in low-level image similarity or differences in the likelihood that the actor was performing the same nameable action? Supplementary analyses indicated that neither the behavioral effects nor the fMRI responses in the MDS and striatum for the prediction task could be accounted for by image similarity or matches in the action being performed (see Supplementary Material). However, when the movie was subsequently restarted, the fMRI response in the MDS and striatal ROIs were significantly

Figure 5. A whole-brain analysis found regions of parietotemporal cortex that were activated when attempting to predict features of a new event. Regions of significant activation are shown, projected on left and right lateral views of the cortical hemispheres using Caret (Van Essen et al., 2001).

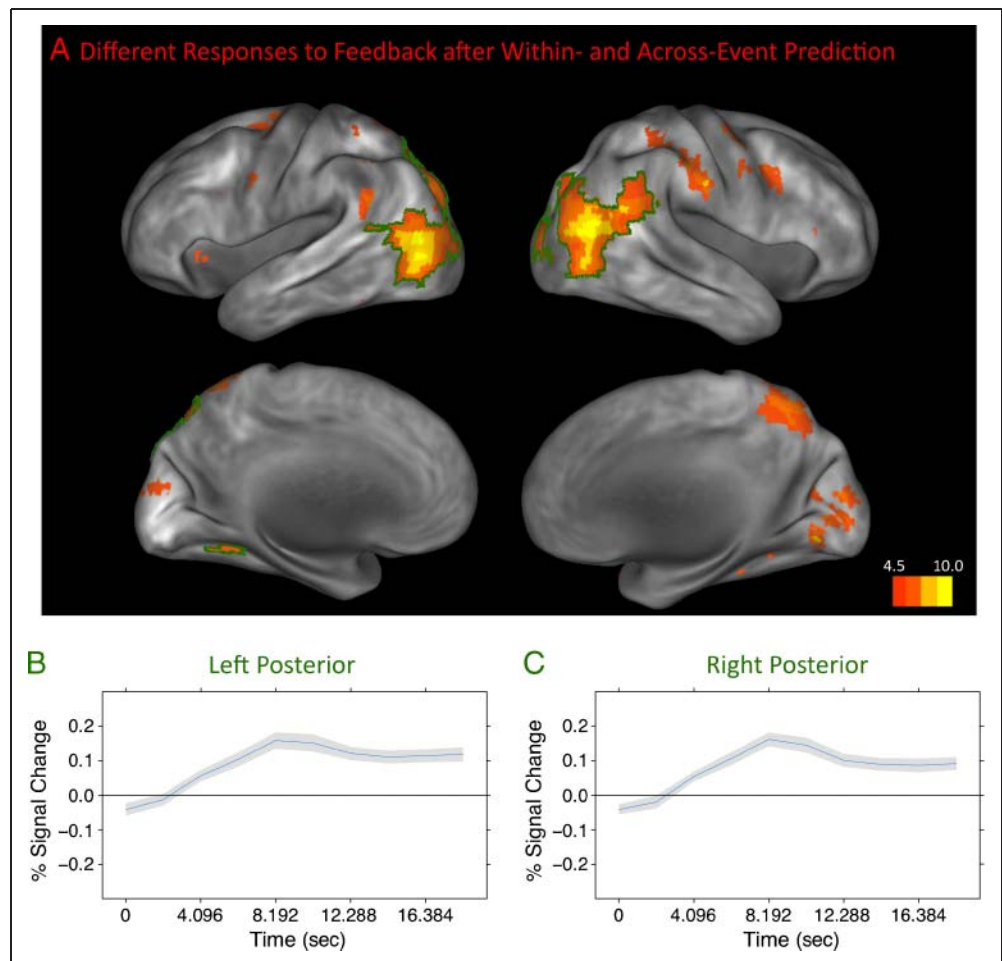


related to image similarity, and controlling for this relationship rendered the event boundary effects nonsignificant. (There were no significant effects of action matches or image similarity in the whole-brain analysis.) Thus, larger responses in the MDS and striatum during this phase may be accounted for by larger visual stimulus changes rather than disconfirming information per se. Because in naturalistic experience prediction errors are likely to be larger when stimulus changes are larger, either mechanism would lead to an adaptive signal for learning and memory updating.

DISCUSSION

In four experiments, human observers attempted to predict the appearance of a visual scene 5 sec in the future of an unfolding naturalistic activity. When the 5-sec interval included an event boundary, predictions were less accurate and were made less confidently. These behavioral markers of unpredictability were accompanied by transient increases in the fMRI response during the attempt to predict in parts of the MDS and in the striatum, which is one of the main projection targets of the MDS. The increases were

Figure 6. Disconfirming information from the beginning of a new event was accompanied by increased fMRI response throughout the cortex. (A) Areas with a significant difference in the time course of response after the onset of the movie in the within- and across-event conditions, projected on the medial and lateral cortical hemispheres using Caret (Van Essen et al., 2001). The two regions that showed the strongest effect are highlighted and plotted in B and C. B and C plot the difference in the fMRI response to restarting the movie in an across-event trial compared with a within-event trial in each region. The new event's beginning was 2.5 sec after the movie onset. Between 8 and 10 sec after the movie onset, a maximal difference response is observed, similar to that observed in the midbrain and striatum (see Figure 4). (Gray regions depict standard errors.)



not accounted for by the degree of visual change in the stimulus or by the occurrence of a change in the action performed by the actor. Thus, they are most consistent with the accumulation of an uncertainty signal mediated by the MDS.

After each prediction attempt the movie resumed, providing the participants with direct feedback regarding the accuracy of their predictions. During this interval, the fMRI signal in other parts of the MDS and striatum increased more for trials that included an event boundary. This response may reflect the prediction error signal associated with receiving disconfirmation (Aron et al., 2004), although it also could reflect greater change in the stimulus.

These results show that the subjective structure of events corresponds with breakdowns in predictability. They are consistent with the proposal that the subjective experience of an event boundary results from transient increases in prediction error (Zacks et al., 2007). However, it is possible that the brain responses observed here reflect some co-occurring attribute of event boundaries other than increased prediction error. We investigated two reasonable candidates: changes in action and changes in image similarity. These control analyses did not provide any evidence that action changes or image similarity were related to effects during the prediction task, which increases our confidence that this response reflects prediction per se. However, it is possible that it could be because of some other uncontrolled feature of the naturalistic stimuli.

The pattern of brain responses observed here converges partially with results from a recent study by Schiffer and Schubotz (2011) that directly manipulated predictability as observers watched sequences of human dance movements after learning to perform similar sequences. In that study, when movement sequences violated viewers' expectations, there were increases in several regions similar to those observed in the current study: the caudate nucleus, the posterior STS and adjacent parietal cortex, medial parietal cortex, and the middle frontal gyrus (see Figures 5 and 6). However, there were also several differences between the responses observed by Schiffer and Schubotz and those seen here. Compared with the present study, they observed responses that were more weighted toward the middle frontal gyrus and the medial pFC and less weighted toward the occipital and posterior/inferior parietal cortex. One possibility is that the greater frontal responses reflect greater engagement of motor simulation because of the dance training, and that the reduced posterior responses reflect greater control over the visual properties of the stimuli. (The SN and VTA were not anatomically identified in that study, so no direct comparisons can be drawn regarding the midbrain responses.)

The results forge a critical link from the statistical violation of expected probabilities as studied in reinforcement learning (Maia, 2009) to the subjective experience of events in time. The activation of the MDS concurrent with prediction failures is consistent with the proposal that prediction

failures engender attentional reorienting and memory updating. Linking attention—and particularly memory updating—to prediction failures may be highly adaptive. If prediction failures tend to occur at those points in time at which a meaningful new event has begun, then that is just the time at which one should update one's mental models of "what is happening now." Disruption of this system may play a role in neuropsychiatric diseases in which dopamine misregulation is implicated, including schizophrenia (Guillin, Abi-Dargham, & Laruelle, 2007) and Parkinson's disease (Olanow & Tatton, 1999).

These results also forge a link from prediction failures observed in simplified laboratory situations, such as in studies of classical conditioning (Schultz, Dayan, & Montague, 1997) and probabilistic classification (Aron et al., 2004), to prediction failures in the comprehension of naturalistic events. They support the hypothesis that dopamine release in the midbrain signals prediction error not only in artificial laboratory paradigms with a discrete "trial" structure and repeated exposures to simplified stimuli, but also during naturalistic events. This mechanism may be profoundly valuable for adaptively regulating immediate behavior as well as for guiding long-term learning about one's environment.

Acknowledgments

This research was supported by NIH grants RO1-MH70674 and T32-AG000030-31 from the National Institutes of Health. The authors thank Alan Anticevic for assistance with FreeSurfer, Thomas Conturo for advice and assistance with MRI pulse sequences, Mark P. McAvoy for assistance with data analysis, Daniel S. Marcus for assistance with data storage and archiving, and Abraham Z. Snyder for advice and assistance with data preprocessing. The authors also thank Sylvia Lee, Joe Dubis, and Albert Deng for assistance with data collection and Deanna Barch, Todd Braver, Jordan Grafman, Joe Magliano, Jesse Sargent, James Zacks, and Rose Zacks for thoughtful comments.

Reprint requests should be sent to Jeffrey M. Zacks, Department of Psychology, Washington University, 1 Brookings Drive, Saint Louis, MO 63130, or via e-mail: jzacks@artsci.wustl.edu.

REFERENCES

- Aron, A. R., Shohamy, D., Clark, J., Myers, C., Gluck, M. A., & Poldrack, R. A. (2004). Human midbrain sensitivity to cognitive feedback and uncertainty during classification learning. *Journal of Neurophysiology*, *92*, 1144–1152.
- Bar, M. (2009). The proactive brain: Memory for predictions. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *364*, 1235–1243.
- Boynton, G. M., Engel, S. A., Glover, G. H., & Heeger, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *Journal of Neuroscience*, *16*, 4207–4221.
- Buckner, R. L., Head, D., Parker, J., Fotenos, A. F., Marcus, D., Morris, J. C., et al. (2004). A unified approach for morphometric and functional data analysis in young, old, and demented adults using automated atlas-based head size normalization: Reliability and validation against manual measurement of total intracranial volume. *Neuroimage*, *23*, 724–738.

- Corbetta, M., Patel, G., & Shulman, G. L. (2008). The reorienting system of the human brain: From environment to theory of mind. *Neuron*, *58*, 306–324.
- D'Ardenne, K., McClure, S., Nystrom, L., & Cohen, J. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science*, *319*, 1264–1267.
- Elman, J. L. (2009). On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon. *Cognitive Science*, *33*, 547–582.
- Enns, J., & Lleras, A. (2008). What's next? New evidence for prediction in human vision. *Trends in Cognitive Sciences*, *12*, 327–333.
- Frith, C. (2007). The social brain? *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *362*, 671–678.
- Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, *66*, 585–595.
- Guillin, O., Abi-Dargham, A., & Laruelle, M. (2007). Neurobiology of dopamine in schizophrenia. *International Review of Neurobiology*, *78*, 1–39.
- Haber, S. N., & Gdowski, M. J. (2004). The basal ganglia. In G. Paxinos & J. K. Mai (Eds.), *The human nervous system* (Vol. 3, pp. 676–738). San Diego, CA: Elsevier/Academic Press.
- Horvitz, J. C. (2000). Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience*, *96*, 651–656.
- Jovicich, J., Czanner, S., Han, X., Salat, D., van der Kouwe, A., Quinn, B., et al. (2009). MRI-derived measurements of human subcortical, ventricular and intracranial brain volumes: Reliability effects of scan sessions, acquisition sequences, data analyses, scanner upgrade, scanner vendors and field strengths. *Neuroimage*, *46*, 177–192.
- Kakade, S., & Dayan, P. (2002). Dopamine: Generalization and bonuses. *Neural Networks*, *15*, 549–559.
- Kurby, C., & Zacks, J. M. (2011). Age differences in the perception of hierarchical structure in events. *Memory & Cognition*, *39*, 75–91.
- Macmillan, N. A., & Creelman, C. D. (1991). *Detection theory: A user's guide*. New York: Cambridge University Press.
- Maia, T. V. (2009). Reinforcement learning, conditioning, and the brain: Successes and challenges. *Cognitive, Affective, & Behavioral Neuroscience*, *9*, 343–364.
- McAvoy, M., Ollinger, J. M., & Buckner, R. L. (2001). Cluster size thresholds for assessment of significant activation in fMRI. *Neuroimage*, *13*, S198.
- McClure, S., Berns, G., & Montague, P. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron*, *38*, 339–346.
- Newton, D. (1976). Foundations of attribution: The perception of ongoing behavior. In J. H. Harvey, W. J. Ickes, & R. F. Kidd (Eds.), *New directions in attribution research* (pp. 223–248). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Niv, Y., & Schoenbaum, G. (2008). Dialogues on prediction errors. *Trends in Cognitive Sciences*, *12*, 265–272.
- Olanow, C., & Tatton, W. (1999). Etiology and pathogenesis of Parkinson's disease. *Annual Review of Neuroscience*, *22*, 123–144.
- Ollinger, J. M., Corbetta, M., & Shulman, G. L. (2001). Separating processes within a trial in event-related functional MRI II: Analysis. *Neuroimage*, *13*, 218–229.
- Pickering, M., & Garrod, S. (2007). Do people use language production to make predictions during comprehension? *Trends in Cognitive Sciences*, *11*, 105–110.
- Reynolds, J. R., Zacks, J. M., & Braver, T. S. (2007). A computational model of event segmentation from perceptual prediction. *Cognitive Science*, *31*, 613–643.
- Schiffer, A.-M., & Schubotz, R. I. (2011). Caudate nucleus signals for breaches of expectation in a movement observation paradigm. *Frontiers in Human Neuroscience*, *5*, 38.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*, 1–27.
- Schultz, W. (2007). Multiple dopamine functions at different time courses. *Annual Review of Neuroscience*, *30*, 259–288.
- Schultz, W. (2010). Dopamine signals for reward value and risk: Basic and recent data. *Behavioral and Brain Functions*, *6*, 24.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.
- Schultz, W., & Dickinson, A. (2000). Neuronal coding of prediction errors. *Annual Review of Neuroscience*, *23*, 473–500.
- Speer, N. K., Reynolds, J. R., & Zacks, J. M. (2007). Human brain activity time-locked to narrative event boundaries. *Psychological Science*, *18*, 449–455.
- Speer, N. K., Swallow, K. M., & Zacks, J. M. (2003). Activation of human motion processing areas during event perception. *Cognitive, Affective & Behavioral Neuroscience*, *3*, 335–345.
- Steward, O. (2000). *Functional neuroscience*. New York: Springer Verlag.
- Summerfield, C., Egner, T., Greene, M., Koechlin, E., Mangels, J., & Hirsch, J. (2006). Predictive codes for forthcoming perception in the frontal cortex. *Science*, *314*, 1311–1314.
- Talairach, J., & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain: 3-Dimensional proportional system, an Approach to Cerebral Imaging*. Stuttgart: G. Thieme.
- Van Essen, D. C., Drury, H. A., Dickson, J., Harwell, J., Hanlon, D., & Anderson, C. H. (2001). An integrated software suite for surface-based analyses of cerebral cortex. *Journal of the American Medical Informatics Association*, *8*, 443–459.
- Volz, K. G., Schubotz, R. I., & von Cramon, D. Y. (2003). Predicting events of varying probability: Uncertainty investigated by fMRI. *Neuroimage*, *19*, 271–280.
- Whitney, C., Huber, W., Klann, J., Weis, S., Krach, S., & Kircher, T. (2009). Neural correlates of narrative shifts during auditory story comprehension. *Neuroimage*, *47*, 360–366.
- Zacks, J. M., Braver, T. S., Sheridan, M. A., Donaldson, D. I., Snyder, A. Z., Ollinger, J. M., et al. (2001). Human brain activity time-locked to perceptual event boundaries. *Nature Neuroscience*, *4*, 651–655.
- Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event perception: A mind/brain perspective. *Psychological Bulletin*, *133*, 273–293.
- Zacks, J. M., Speer, N. K., Vettel, J. M., & Jacoby, L. L. (2006). Event understanding and memory in healthy aging and dementia of the Alzheimer type. *Psychology & Aging*, *21*, 466–482.
- Zink, C. F., Pagnoni, G., Chappelow, J., Martin-Skurski, M., & Berns, G. S. (2006). Human striatal activation reflects degree of stimulus saliency. *Neuroimage*, *29*, 977–983.
- Zink, C. F., Pagnoni, G., Martin, M. E., Dhamala, M., & Berns, G. S. (2003). Human striatal response to salient nonrewarding stimuli. *The Journal of Neuroscience*, *23*, 8092–8097.